A Canon Communications LLC Publication

EDN

VOICE OF THE ENGINEER

JUNE 10
Issue 11/2010
www.edn.com

ALTERING THE SOC-DESIGN FLOW

Page 30

United for a better world

NEC Electronics Corporation and Renesas Technology Corp. have combined their businesses to launch

Renesas Electronics Corporation.

Renesas Electronics has integrated the microcontroller, power and analog semiconductors, as well as the advanced technologies and strong experience of both companies, to provide system solutions for the global market. As the #1 supplier of microcontrollers in the world*, the new Renesas Electronics has a corporate vision to harness the power of semiconductors to realize intelligent devices that can enable a sustainable planet and improve the lives of both current and future generations. Together we can realize a society where life can be enjoyed in harmony with the planet.

* Source: Gartner, "Semiconductor Applications Worldwide Annual Market Share: Database" 25 March 2010
This is the 2009 ranking based on revenue.

www.renesas.com

Renesas Electronics

RENESAS

COUNTLESS OPTIONS!

SAME-DAY SHIPPING on orders entered by 8:00 pm CENTRAL
NEXT-DAY DELIVERY

More Than 465,000 Products in Stock

Over 45,000 NEW Products Added in the Last 90 DAYS

440+ Suppliers*

Digi-Key® CORPORATION

#1 Breadth of Product Line
#1 Availability of Product
#1 Overall Performance

Source: EETimes Distribution Study, August 2009

The industry's broadest product selection available for immediate delivery

www.digikey.com
1.800.344.4539

*Digi-Key is an authorized distributor for all supplier partners. New products added daily. © 2010 Digi-Key Corporation, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA

NEDA Supplier Authorized Distributors

Advanced Portable System Power Management ICs

Quad Output PMIC with HyperLight Load™ DC/DC, 3 LDOs and I²C Control

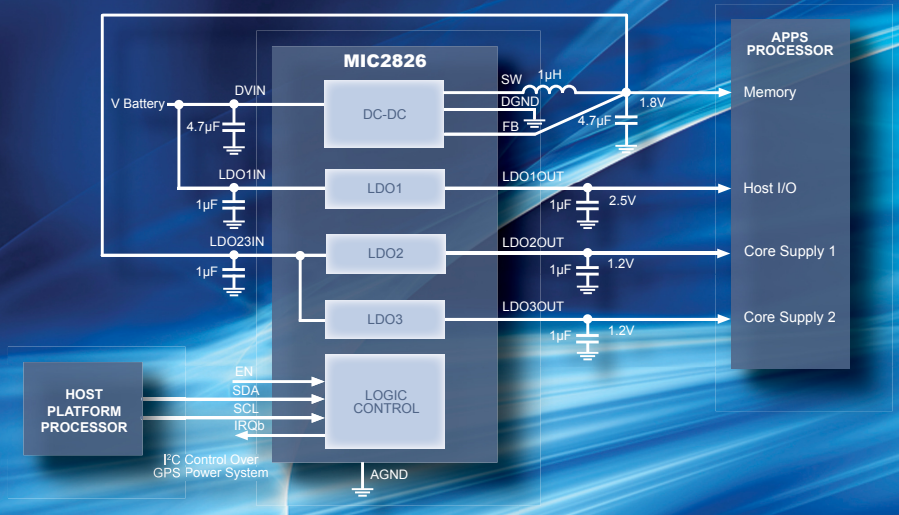

- Fast-mode 400kHz I²C Control Interface
- Dynamic Voltage Scaling & Sequencing
- Extremely High Light Load Efficiency
- Tiny 2.5mm x 2.5mm MLF® Package

Today's advanced portable electronics require highly integrated power system solutions for complex system power, small size, and long battery life. Micrel provides complete system power solutions, such as the MIC2826, to meet these needs.

The MIC2826 integrates a single 500mA HyperLight Load™ synchronous buck (step-down) regulator with three Low Dropout Regulators in a tiny 2.5mm x 2.5mm Thin MLF® package. The MIC2826 also features a standard 400kHz Fast-mode I²C interface that provides Dynamic Voltage Scaling (DVS), programmable power sequencing, and individual output control, making it the perfect integrated power system solution.

For more information, contact your local Micrel sales representative or visit Micrel at: www.micrel.com/ad/mic2826.

Ideal for use in:

- Application Processors
- GPS Subsystems
- Mobile Phones/PDAs
- Portable Media Players
- General Purpose PMIC

MICREL®
Innovation Through Technology®

www.micrel.com

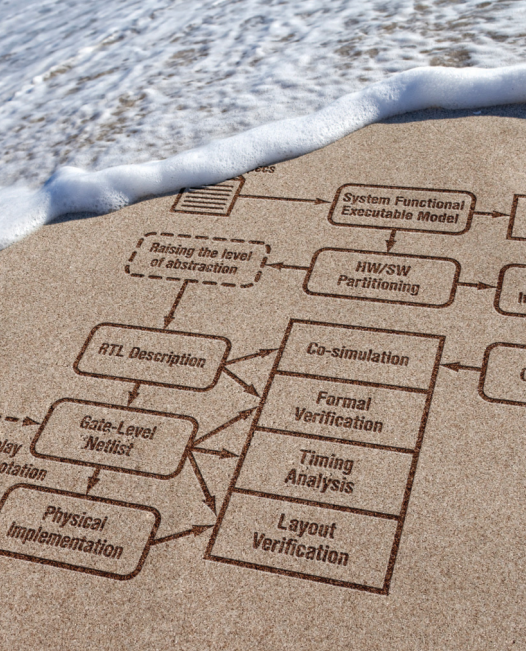

EDN 6.10.10 contents

Altering the SOC-design flow

By Ron Wilson, Executive Editor

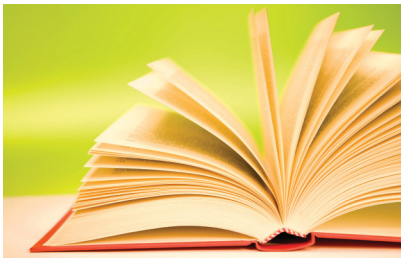

Low power: a chip and system-design primer

By Mike Salas, Silicon Labs

Power-grid analysis on SOC-graphics-chip design

By Ravi Poddar, PhD, and Ted Sun, Integrated Device Technology Inc

Test your testers for R&R

By Martin Rowe, Senior Technical Editor, Test & Measurement World

pulse

User-configurable FPGA modules boost PXI-system versatility

By Richard A Quinnell, Contributing Technical Editor

DESIGNIDEAS

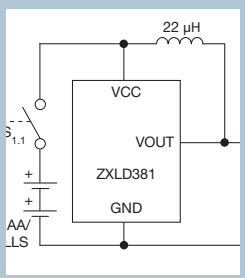

“I need a faster, more accurate audio analyzer.”

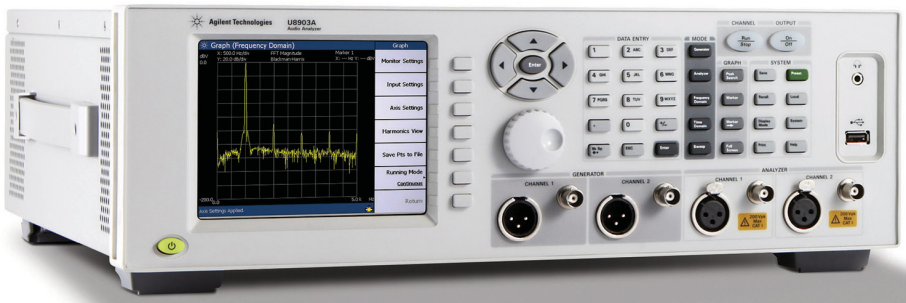

We knew you’d say that.

The new Agilent U8903A audio analyzer is four times more accurate and over seven times faster than its popular HP 8903B ancestor. In fact, it is among the fastest and most accurate currently in its class. It’s also a more fully capable audio analyzer, letting you make crosstalk measurements and phase measurements, FFT analysis, and graph sweep. That’s taking it to the next level. That’s Agilent.

Learn how to migrate effortlessly -
Free application note filled with tips
www.agilent.com/find/AgilentU8903A

u.s. 1-800-829-4444 canada: 1-877-894-4414

Agilent Technologies

contents 6.10.10

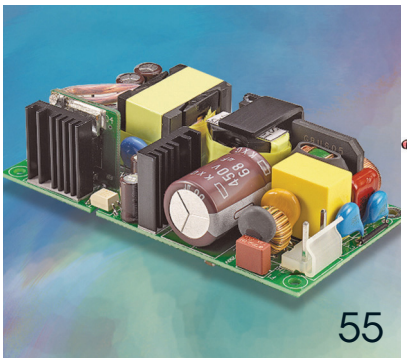

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ENGINEERS' TRUE STORIES

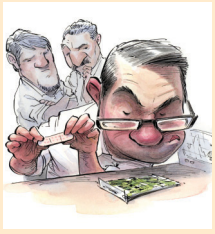

In *EDN*'s Tales from the Cube, engineers relate their most vexing design challenges—and how they conquered them. Here are some examples of how your peers solved real-world problems:

The turn of the screwdriver
➔www.edn.com/100610toca

Sting like a bee
➔www.edn.com/100610tocb

Can't put a Band-Aid on a boomerang
➔www.edn.com/100610tocc

Seeing red over dead LEDs
➔www.edn.com/100610tocf

Read more in the Tales archive at www.edn.com/tales.

FROM *EDN*'s BLOGS

Desolder SMD ICs with bismuth solder
From Anablog, by Paul Rako
When I wrote my article on prototyping, I learned that one way to desolder chips and such is to melt some bismuth solder into the solder joints.
➔www.edn.com/100610tocd

Smart-grid chips integrate functions as meter vendors aim for standards
From PowerSource, by Margery Conner
Freescale has fleshed out its smart-power-meter-IC family with the announcement of the MC9S08GW64, which integrates gas- and water-metering functions. The chip, based on an 8-bit S08 core, includes an electricity-metering analog front end with two independent 16-bit SAR ADCs and a programmable delay block for phase-error compensation.
➔www.edn.com/100610toce

EDN® (ISSN#0012-7515) is published semimonthly, 24 times per year, by Canon Communications LLC, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 310/445-4200; FAX 310/445-4299. Periodicals postage paid at Los Angeles, California, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. Rates for nonqualified subscriptions, including all issues: US, $150 one year; $250 two years; $300 three years. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. For telephone inquiries regarding subscriptions, call 763/746-2792. Email: EDN@kmpsgroup.com. CHANGE OF ADDRESS—Notices should be sent promptly to PO Box 47461, Plymouth, MN 55447. Please provide old mailing labels as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publisher cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 47461, Plymouth, MN 55447. Canada Post: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to Pitney Bowes, Inc., PO Box 25542, London, ON N6C 6B2.

DENSE. FAST. FLEXIBLE.

SEARAY™
OPEN PIN FIELD INTERCONNECTS

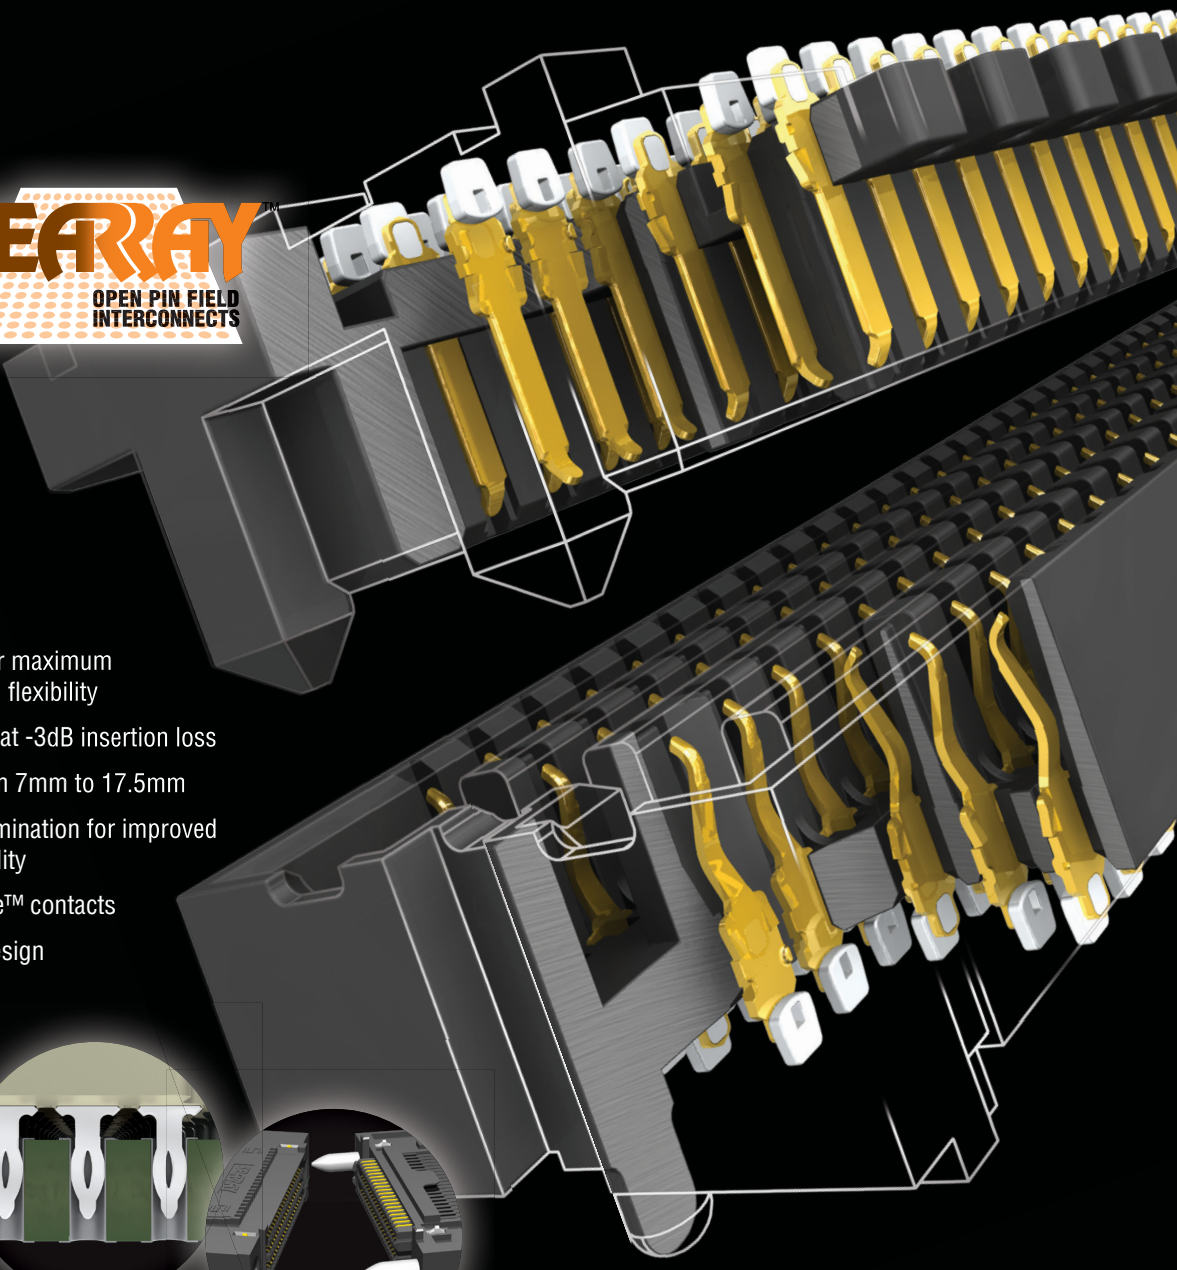

- Up to 500 I/O's for maximum grounding/routing flexibility
- Rated at 9.5 GHz, at -3dB insertion loss
- Stack heights from 7mm to 17.5mm
- Solder charge termination for improved processing reliability
- Rugged Edge Rate™ contacts
- 0,8mm pitch in design

Press-Fit Option

Right Angle Version

samtec

Transmission Line Solutions

www.samtec.com/searay

BY RON WILSON, EXECUTIVE EDITOR

The crash of 2:45 pm, network theory, and the future of SOC design

As the May 6 New York Stock Exchange trading event passes from hard news into the hands of the television experts, most fingers are pointing at an erroneous trade that suddenly dumped an absurd number of shares for sale onto the market, causing the loss of more than 1000 points on the Dow in less than a half-hour. Computer programs order most of the NYSE trades, however, and humans never see them, so I think we need to consider another line of inquiry.

The data centers that run these algorithmic trading programs connect to even bigger data centers that log, execute, and report the trades. These connections are low-latency because a change of a few microseconds in the round-trip latency can change the income or loss an algorithm generates at a node. This loss can be enough to drive banks to pay a lot for slightly faster hardware. Further, the data centers can connect to places other than the trading hub, including to each other.

This situation represents a mesh network in which the topology is undocumented and dynamic. The variability in link latencies is sufficient to influence the behavior of the nodes. Except for the exchanges at the hubs, the algorithm running at each node is secret from every other node, and even its authors may not fully understand it. In other words, these data centers contain a recipe for a dynamically unstable network that no one can analyze.

These networks have few stability mechanisms. One is the hypothetical tendency for free markets to discover the current fair value of a commodity. In principle, this tendency should lead the networks to converge on "fair" prices for stocks. Analysts have not studied this hypothesis in cases in which the agents in the market are algorithms that should outsmart each other. The other stability mechanism is the willingness of humans to put an artificial brake on trading speed when they become frightened. But such a decision takes tens of minutes.

In this incomprehensible, unstable network with only weak and glacially slow controls, you would expect just the sort of behaviors we are starting to see: local oscillations and limit cycles and the occasional global instability.

According to Carnegie Mellon University Professor Jeannette Wing, assistant director for CISE (computer and information science and engineering) at the National Science Foundation, the world faces an analogous problem with global information networks of all sorts: We just don't know how they behave. One of the challenges most interesting to Wing's branch of the NSF is to see the construction of a theory of large networks so we can begin to understand them.

What does this information have to do with SOCs (systems on chips)? As we move deeper into heterogeneous multicore SOCs, we are beginning to build large on-chip networks that have mesh organization and nondeterministic latencies. They also have algorithmic nodes. Only the IP (intellectual-property) provider that created these nodes may know about their actual operation. These nodes will compete for shared memory, interconnect quality of service, and—in the future—power.

Without a useful theory of network behavior at the chip level, we may be unable to model SOCs well enough to predict what they will do in a system.

Without a useful theory of network behavior at the chip level, we may be unable to model SOCs well enough to predict what they will do in a customer's system. Unless we develop such a theory for financial networks, however, our ability to design SOCs won't much matter. EDN

Contact me at ronald.wilson@cancom.com.

Support **Across The Board.**™
From Design to Delivery™

Now, you can have it all.™

Faster and easier than ever before. Our commitment to customer service is backed by an extensive product offering combined with our supply chain and design chain services – which can swiftly be tailored to meet your exact needs. We have dedicated employees who have the experience to provide the highest level of customer service with accuracy and efficiency. All of our technical experts are factory certified on the latest technologies, providing you the expertise to move projects forward with speed and confidence.

Avnet offers the best of both worlds: extensive product and supply chain knowledge, and specialized technical skill which translates into faster time to market – and the peace of mind that comes from working with the industry's best. **Avnet is ranked Best-In-Class* for well-informed sales reps, knowledgeable application engineers and our design engineering services – proof that we consistently deliver:**

> Industry recognized product expertise
> Specialized technical skills

Ready. Set. Go to Market.™
Visit the Avnet Design Resource Center™ at:
www.em.avnet.com/drc

AVNET®
electronics marketing

Accelerating Your Success™

Avnet Green Initiative

Supplier Authorized Distributor

1 800 332 8638
www.em.avnet.com

*As rated by Hearst Electronics Group: The Engineer & Supplier Interface Study, 2009.

Follow us on Twitter!
www.twitter.com/avnetdesignwire

EDN

PUBLISHER, *EDN* WORLDWIDE
Russell E Pratt, 1-781-869-7982;
russell.pratt@cancom.com

ASSOCIATE PUBLISHER, *EDN WORLDWIDE*
Judy Hayes, 1-925-736-7617;
judy.hayes@cancom.com

EDITOR-IN-CHIEF, *EDN WORLDWIDE*
Rick Nelson, 1-781-869-7970;
richard.nelson@cancom.com

EXECUTIVE EDITOR
Ron Wilson, 1-510-744-1263;
ronald.wilson@cancom.com

MANAGING EDITOR
Amy Norcross
1-781-869-7971;
fax: 1-781-862-4853;
amy.norcross@cancom.com
Contact for contributed technical articles

ANALOG
Paul Rako, Technical Editor
1-408-745-1994;
paul.rako@cancom.com

MASS STORAGE, MULTIMEDIA, PCs, AND PERIPHERALS
Brian Dipert, Senior Technical Editor
1-916-760-0159;
brian.dipert@cancom.com

NEWS
Suzanne Deffree, Managing Editor
1-631-266-3433;
suzanne.deffree@cancom.com

POWER SOURCES, ONLINE INITIATIVES
Margery Conner, Technical Editor
1-805-461-8242;
fax: 1-805-461-9640;
margery.conner@cancom.com

DESIGN IDEAS EDITOR
Martin Rowe,
Senior Technical Editor,
Test & Measurement World
edndesignideas@cancom.com

SENIOR ASSOCIATE EDITOR
Frances T Granville
1-781-734-8439;
fax: 1-781-862-4853;
f.granville@reedbusiness.com

CONSULTING EDITOR
Jim Williams, Staff Scientist,
Linear Technology
edn.editor@cancom.com

CONTRIBUTING TECHNICAL EDITORS
Dan Strassberg,
strassbergedn@att.net
Robert Cravotta,
robert.cravotta@embeddedinsights.com

COLUMNISTS
Howard Johnson, PhD, Signal Consulting
Bonnie Baker, Texas Instruments
Pallab Chatterjee, SiliconMap
Kevin C Craig, PhD, Marquette University

LEAD ART DIRECTOR
Marco Aguilera

ASSOCIATE ART DIRECTOR
Tim Burns

PRODUCTION
Michael Ciardiello,
Director of Premedia Technologies
Jeff Tade,
Production Director
Brian Wu,
Publications Production Manager
Jeff Polman, Derric Treece,
Senior Production Artists
William Baughman, Ricardo Esparza,
Production Artists

EDN EUROPE
Graham Prophet,
Editor, Reed Publishing
+44 118 935 1650;
gprophet@reedbusiness.fr

EDN ASIA
Luke Rattigan,
Chief Executive Officer
luke.rattigan@rbi-asia.com
Kirtimaya Varma,
Editor-in-Chief
kirti.varma@rbi-asia.com

EDN CHINA
William Zhang,
Publisher and Editorial Director
wmzhang@rbichina.com.cn
Jeff Lu,
Executive Editor
jefflu@rbichina.com.cn

EDN JAPAN
Katsuya Watanabe,
Publisher
k.watanabe@reedbusiness.jp
Ken Amemoto,
Editor-in-Chief
amemoto@reedbusiness.jp

EXECUTIVE OFFICERS
Charles G. McCurdy,
Chief Executive Officer
Fred Gysi,
Chief Financial Officer
Mike Deering,
Chief Technology Officer
Ron Wall,
Senior Vice President, Publications
Kevin O'Keefe,
Senior Vice President, Events Division
Roger Burg,
Vice President, Operations
Jason Brown,
Vice President, E-media

EDN. 33 Hayden Avenue, Lexington, MA 02421. **www.edn.com**. **Subscription inquiries:** 1-763-746-2792; EDN@kmpsgroup.com. **Address changes:** Send notice promptly to PO Box 47461, Plymouth, MN 55447. Please provide an old mailing label as well as your new address. Allow two months for the change. Canon Communications LLC, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 1-310-445-4200; fax: 1-310-445-4299.

CANON COMMUNICATIONS LLC

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

JFET op amps suit use in professional-audio applications

Texas Instruments' Burr-Brown Division recently released the single OPA1641, dual OPA1642, and quad OPA1644 operational amplifiers, which feature a JFET input stage. The parts operate with ±2.25 to ±18V power supplies, and output distortion is 0.00005% at 1 kHz. The devices feature a ±20-pA maximum bias current at 25°C and a 70-pA typical bias current at 85°C. The 11-MHz-bandwidth IC has less than 3.5-mV maximum offset voltage. It consumes 1.8 mA quiescent current per channel and achieves a 20V/µsec slew rate. The amp is unity-gain stable and has an open-loop gain of 105 dB. The rail-to-rail output has a 2-kΩ load. The amplifiers have 5.1-nV/√Hz voltage noise at 1 kHz and 8 nV/√Hz at 10 Hz.

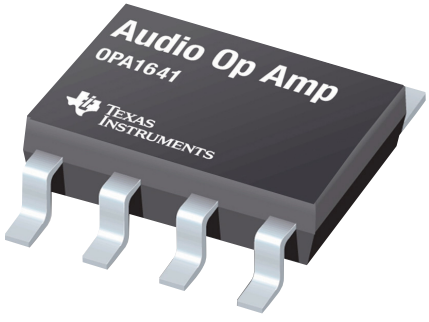

The OPA1641 JFET op amp features low noise and small input bias currents, suiting use as an amplifier on the outputs of an audio codec.

Because JFETs are buried devices, they are not subject to noise that surface defects create. These amplifiers' JFETs also feature a flicker-noise corner of less than 20 Hz, making them ideal for professional-audio applications. The low bias-current inputs also suit them for use with signals having high source impedance.

The single-channel OPA1641 is available in an SO-8 package for 95 cents (1000). The two-channel OPA1642 is available in an SO-8 or an MSOP-8 package for $1.45, and the four-channel OPA1644 is available in an SO-14 or a TSSOP-14 package for $1.95.

—by Paul Rako

▷**Texas Instruments**, www.ti.com.

FEEDBACK LOOP

"As I understand 'The Bumblebee Myth,' the bumblebee was proved to have insufficient wing area to enable it to fly. However, the equations used were the ones used for fixed-wing aircraft! And the bumblebee's wings do, indeed, move!"

—Computer geek and *EDN* reader Paul E Musselman, in *EDN*'s Feedback Loop, at www.edn.com/article/457778-Sting_like_a_bee.php. Add your comments.

20-bit DAC simplifies MRI-signal path

Analog Devices' new AD5791 DAC has better performance and uses fewer parts in MRI (magnetic-resonance-imaging) equipment than other methods targeting this application. It also finds use in process control; data-acquisition systems; source-measurement units; and general analog applications, such as digital gain and offset adjustment, programmable-voltage and -current sources, and programmable attenuators.

The device offers a maximum-relative-accuracy specification of ±1 LSB, guaranteed monotonic operation with a ±1-LSB-maximum DNL (differential-nonlinearity) specification, 0.025-ppm low-frequency noise, and 0.05-ppm/°C output drift. The spectral-noise density is 9 nV/√Hz, and settling time is 1 µsec. You can configure the output for standard unipolar 5 and 10V or bipolar ±5 and ±10V ranges. The device operates from a 33V bipolar power supply and consumes 4.2 mA of current on the analog power pins and 800 mA on the digital power pin. It directly drives a 60-kΩ load, eliminating linearity and noise degradation that are the drawbacks of output-buffering op amps. The AD5791 uses a three-wire serial interface that operates at clock rates as high as 50 MHz.

The part comes in 24-pin LFCSPs and 20-pin TTSOPs, operates over a –40 to +105°C temperature range, and has a suggested retail price of $38 (1000). Samples are available now, and production quantities will become available in August 2010.

—by Paul Rako

▷**Analog Devices**, www.analog.com.

MEET THE GUY that ELIMINATED HIS TEAM'S MANUFACTURING VARIABILITY ISSUES.

FIX YOUR MANUFACTURING VARIABILITY PROBLEMS AND YOUR REPUTATION WILL PRECEDE YOU. If you're designing chips for high functionality, high speed and lower power consumption at the most advanced process nodes, you've got variability issues. We have the solution that will increase your yield, performance and prestige by a wide margin. | Get more information at mentor.com/solutions/manufacturing-variability.

Mentor Graphics®

THE EDA TECHNOLOGY LEADER

Adlink introduces Atom-based SBCs

Adlink Technologies has introduced two SBCs (single-board computers) that it based on versions of Intel's (www.intel.com) Atom processor. The new EPIC (embedded platform for industrial computing) SBC, the Ampro ReadyBoard 740, integrates a dual-core Atom D510 or a single-core Atom D410 processor; the Ampro CoreModule 740SFF (small-form-factor) PC/104-Plus uses the Atom N450 processor.

The ReadyBoard 740's 1.66-GHz processor includes an Intel ICH8M chip set, an onboard solid-state drive, an H.264 hardware video decoder, networking capability, and I/O functions. One SO-DIMM accommodates as much as 2 Gbytes of 667-MHz DDR2 RAM. The board targets use in applications in harsh environments that require high-end video performance, such as vehicle computing platforms, self-service kiosks, digital signage, and video surveillance.

Adlink's optional MiniModule PWR power-protection module adapts the board to meet the demands of in-vehicle power environments. MiniModule PWR accommodates wide input-voltage ranges and offers reverse-battery and overvoltage protection. The board has an operating-temperature range of –20 to +70°C; prices start at $410.

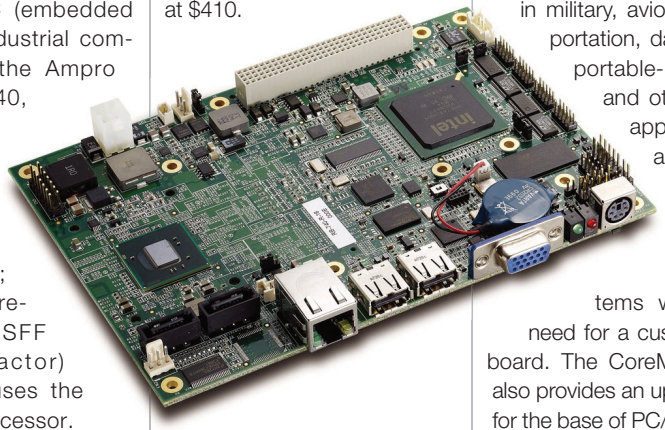

The ReadyBoard 740's 1.66-GHz processor includes an Intel ICH8M chip set, an onboard solid-state drive, an H.264 hardware video decoder, networking capability, and I/O functions.

The CoreModule 740 utilizes the Atom's two-chip architecture with integrated memory to provide a balance of power consumption and performance. The module is suitable for use in sealed enclosures and combines legacy I/O interfaces, including ISA (industry-standard architecture), IDE (integrated device electronics), and serial and parallel ports with onboard video, 1 Gbyte of soldered RAM, and USB (Universal Serial Bus) 2.0 and the PCI (Peripheral Component Interconnect) bus, in a 90×96-mm footprint.

The PC/104-Plus stackable form factor allows OEMs in military, avionics, transportation, data-logging, portable-computing, and other rugged applications to add an Intel Architecture controller to their systems without the need for a custom carrier board. The CoreModule 740 also provides an upgrade path for the base of PC/104 system designs. Prices start at $400 (OEM quantities).

Adlink also announced a program to assist companies in migrating from products affected by Intel's recent product-change notifications, announcing the discontinuance of key embedded processors and chip sets.

Adlink will help its customers extend product life cycles, review customer application-system architectures, provide recommendations for upgrade paths, provide gap analysis, and offer level drop-in replacement boards.

—by Rick Nelson

▷**Adlink Technologies**, www.adlinktech.com.

Dilbert By Scott Adams

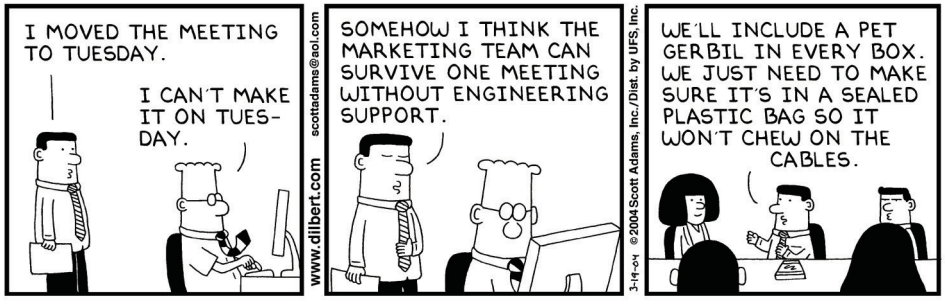

DIGITAL PFC IC UPS PERFORMANCE, SHRINKS PARTS COUNT AND PRICE

06.10.10

Digital chips continue to encroach on traditionally analog parts of power management. For example, Cirrus Logic recently announced what it claims is the first integrated, high-volume IC for digitally controlled PFC (power-factor correction). The DCM (discontinuous-conduction-mode), active-PFC CS1500 and CS1600 ICs target power supplies that operate at 300W or less. The CS1500 addresses power supplies for applications such as laptops, digital TVs, and PCs, and the CS1600 targets electronic lighting ballasts.

At prices of approximately 30 cents each (high volumes), the chips compare in price with analog PFC ICs but require 30% fewer additional components and parts for EMI (electromagnetic-interference) filtering. The power factor, which varies with the input line voltage and with load, is greater than 0.95.

—by Margery Conner

▷**Cirrus Logic**, www.cirrus.com.

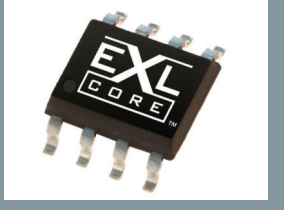

Weighing less than 6.2 lbs, the FieldFox 4- and 6-GHz, two-port VNAs boast a novel internal-calibration system.

The Power Behind the Green

High-efficiency, high-density power conversion

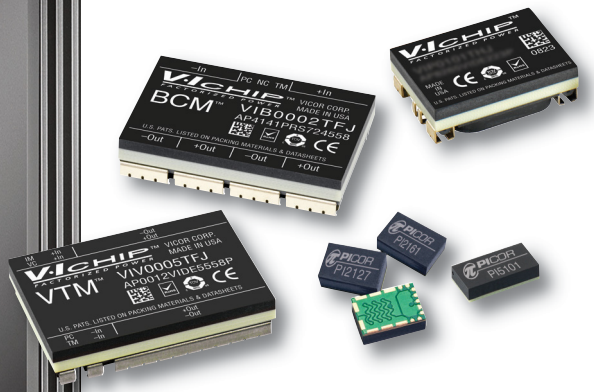

Vicor's powertrain and power management technologies enable extreme computing performance and efficiency.

Three of the top 10 supercomputers in the TOP500 list (TFLOPS) and twelve of the top 20 Green500 list (MFLOPS/Watt) are powered by Vicor.

Data centers using Vicor technology save $600k in electricity, 2,700 tons of CO_2, and 2,500 barrels of oil per year.

System features include:

- High-efficiency, high-voltage power distribution
- Direct conversion to processors, memory and drives
- Power conversion > 1000 W/in^3
- Highly efficient, very small size fault protection solutions for redundant power systems

Call 800-735-6200 to learn more about the power of Vicor technology.

V·I CHIP FACTORIZED POWER

vicorpower.com/datacenter/edn

VICOR

LeCroy debuts embedded-test tools

LeCroy recently introduced the ArbStudio AWGs (arbitrary-waveform generators) and LogicStudio 16 logic analyzer at the Embedded Systems Conference (www.embedded.com), which took place in April in San Jose, CA. The ArbStudio AWGs generate signals as fast as 125 MHz and include PWM (pulse-width-modulation) capabilities. The software interface that controls the hardware simplifies waveform creation with a navigation tree that allows easy access to all channels.

The ArbStudio series includes four models: two- and four-channel versions with analog-waveform capabilities plus two- and four-channel versions offering a combination of analog-waveform- and digital-pattern-generation capabilities. The four-channel models have an expansion port that allows you to connect as many as eight units. All models have a 125-MHz bandwidth, 1G-sample/sec maximum sample rate, 2M-point/channel memory, and 16-bit resolution. The instruments support both true arbitrary and DDS (direct-digital-synthesis) technologies. ArbStudio software runs on an external PC. The base price for ArbStudio ranges from $2490 to $4990.

ArbStudio AWGs generate signals as fast as 125 MHz and include PWM capabilities.

The LogicStudio 16 brings logic-analyzer functions to a PC, providing 16 channels with a sample rate of 1G sample/sec and maximum input frequency as high as 100 MHz. LogicStudio 16 software provides a dynamic waveform display with an intuitive user interface.

Tools for digital debugging include timing cursors, zooming and panning functions, a persistence display, and a history mode that can replay old data captures. LogicStudio supports protocol analysis for I²C (inter-integrated-circuit) interfaces, SPIs (serial-peripheral interfaces), and UART (universal asynchronous-receiver/transmitter) interfaces. It can trigger on bus addresses or data packets.

LogicStudio provides a communication link to LeCroy's WaveJet oscilloscope, thereby turning a PC into a mixed-signal debugging environment. The base price for LogicStudio 16 is $990. **—by Rick Nelson**

▷**LeCroy**, www.lecroy.com.

Three-phase, step-down dc/dc controller targets high current rails

Linear Technology's new three-phase, single-output LTC3829 synchronous step-down dc/dc controller features efficiency as high as 95%, PolyPhase operation, differential output-voltage sensing, and integrated PLL (phase-locked-loop) clock synchronization. You can parallel as many as six phases and clock them out of phase to minimize input- and output-filtering requirements for applications requiring currents as high as 150A.

The differential amplifier provides true remote output-voltage sensing of both the positive and the negative terminals, enabling high-accuracy regulation independent of IR losses in trace runs, vias, and interconnects. Applications include high-current ASIC and FPGA supplies, power-distribution buses, and high-power audio amplifiers.

The LTC3829 operates with all N-channel MOSFETs from input voltages of 4.5 to 38V, and it can produce ±0.75%-accurate output voltages of 0.6 to 5V. You can ascertain the output current by monitoring the voltage drop across the output inductor's DCR (dc resistance) or by using a sense resistor. Programmable DCR temperature compensation maintains an accurate and constant current-limit setpoint over a broad temperature range. Onboard gate drivers minimize MOSFET switching losses and allow the use of multiple MOSFETs that connect in parallel. You can program a fixed operating frequency of 250 to 770 kHz or synchronize the frequency to an external clock with the internal PLL. The device's minimum on-time of 90 nsec makes it ideal for high-step-down-ratio applications.

The LTC3829 incorporates an adjustable "stage-shedding" technique to increase light-load efficiency by eliminating the gate-charge and switching losses of two of its output stages. You can also configure the LTC3829 controller for adjustable burst-mode operation, producing higher efficiency at light loads. A nonlinear control mode is optional and improves load-step transient response. Adaptive voltage positioning minimizes the maximum transient voltage deviation during a step load.

Tracking and sequencing functions allow the optimization of power-up and power-down of multiple power supplies. Additional features include current-mode control, an onboard low-dropout regulator for IC power,

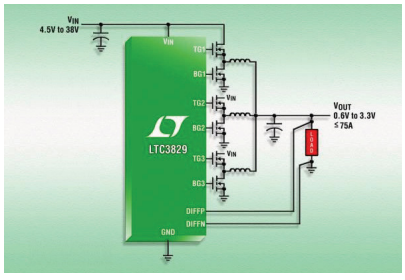

The three-phase, single-output LTC3829 synchronous step-down dc/dc controller features efficiency as high as 95%.

programmable soft start, a power-good output, and an external power supply.

The LTC3829 is available in thermally enhanced, 38-lead TSSOP or 38-lead, 5×7-mm QFN packages and operates over a –40 to +125°C operating junction-temperature range, with prices starting at $3.71 (1000). **—by Fran Granville**

▷**Linear Technology**, www.linear.com.

Compact high power

Industrial or Medical

SHP Series

MHP Series

650 W

- 650 & 1000 Watt Models
- Rugged Industrial Construction
- Medical Safety Approvals (MHP range)
- Top or End Fan & 'U' Channel Options on 650W
- -20 ºC to +70 ºC Operation
- Variable Fan Speed for Noise Reduction
- 5 V Standby
- AC OK, Remote On/Off & Active Current Share
- 3 Year Warranty

1000 W

Call or visit our website to request a copy of our Power Supply Guide 09/10 and see our complete line of power products.

Toll free: 1-800-253-0490
nasales@xppower.com

Also available from:
Newark
www.newark.com

XP Power
www.xppower.com

Backplanes target embedded-system designs

At the ESC (Embedded Systems Conference, www.embedded.com) held in April in San Jose, CA, National Instruments took aim at helping designers integrate PXI (Peripheral Component Interconnect extensions for instrumentation) and CompactPCI (Peripheral Component Interconnect) instruments into their embedded systems by introducing board-level backplanes. The backplanes, previously available only as part of NI's PXI/CompactPCI and PXIe (PXI Express) chassis, allow OEMs to create their own custom, rugged enclosures that can accommodate PXI, PXIe, CompactPCI, and CompactPCIe modules.

New backplanes from National Instruments allow OEMs to create their own custom, rugged enclosures that can accommodate PXI, PXI Express, CompactPCI, and CompactPCIe modules.

The more than 10 new 3U and 6U backplanes offer four to 18 slots. Engineers can design custom installations and enclosures around the backplanes and integrate more than 1500 PXI modules, including data-acquisition cards; FPGA-based I/O modules; high-end instruments, such as signal generators and RF-signal analyzers; and a variety of bus-interface modules, including serial, MIL-STD (military standard)-1553, IEEE 1588, Profibus, and DeviceNet versions.

Designers are able to use the NI LabView graphical-system-design platform to design, prototype, and deploy all aspects of their systems, in keeping with a focus on letting domain experts in robotics, medical, and energy industries, for example, play a significant role in embedded-system designs, says Casey Weltzin, NI's LabView real-time product manager.

—by Rick Nelson

▷**National Instruments**, www.ni.com.

PMC goes multicore with RAID controller

Multicore SOC (system-on-chip) architectures are gradually seeking out the spectrum in which no other way exists to reach the desired performance level, but also in midrange applications, in which several small cores can be cheaper and more energy-efficient than one fire-breathing CPU. Consequently, embedded-application developers are wrestling their code into multiple threads, and SOC developers are providing cores to match the threads.

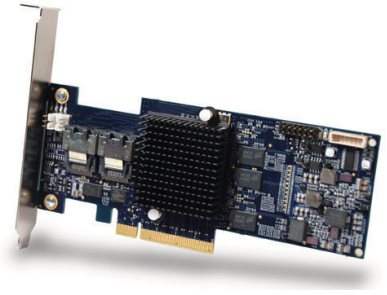

PMC-Sierra's new line of RAID-controller ICs, a joint development with IBM (www.ibm.com), features a multithreaded version of IBM's RAID firmware stack and a low-profile, eight-lane PCIe Generation 2 card to plug into x86 servers.

A case in point, as PMC-Sierra recently announced, is a RAID (redundant-array-of-inexpensive-disks) controller. This application might not seem computationally intensive, but, as Cameron Brett, PMC-Sierra's product-marketing manager, points out, the plummeting price of NAND flash is changing the picture.

Consider a high-performance RAID of, say, 20 rotating drives. If you want performance, you will use the drives in short-stroke mode—that is, employing only the outermost tracks, on which the data rate is highest and you can minimize seek time. That approach gives you good performance, but you use less than 10% of each drive's raw capacity. For about the same price, you could get four 100-Gbyte solid-state drives, connect them in a RAID configuration, and get about 300 Gbytes of usable capacity. The solid-state drives would slash your power consumption from more than 200W to about 10W, however, and they would increase available throughput from about 9000 to approximately 90,000 IOPS (input/output operations per second). Suddenly, the performance of the RAID controller is an issue.

PMC has responded with a new line of RAID-controller ICs, a joint development with IBM (www.ibm.com) to produce a multithreaded version of IBM's RAID firmware stack, and a low-profile, eight-lane PCIe (Peripheral Component Interconnect Express) Generation 2 card to plug into x86 servers. The products target use in RAID configurations of solid-state drives, as well as conventional hard-disk-drive RAID arrays, using either SATA (serial advanced-technology attachment) or SAS (serial attached small-computer-system interface), including the 6-Gbps SAS generation, so the controller must deliver performance. "Each solid-state drive is capable of 15,000 to 30,000 IOPS," says Zaki Hassan, PMC's director of product marketing. Even with a four-drive array, this rate is significant.

To support the RAID stack at this throughput, PMC clusters three multithreaded MIPS (www.mips.com) 34k cores around a high-speed switch. The chip uses only local memory without external DRAM, Brett says. The board that uses the chip, the BR5225-80, measures in at more than 136,000 IOPS on 4000 random reads, and more than 44,000 IOPS on the 4K OLTP (online-transaction-processing) random 2-to-1 metric. Additionally, the chip supports one end of a Web-accessible, SMI-S (Storage Management Initiative Specification) 1.4-compliant management utility that runs on the host server.

—by Ron Wilson

▷**PMC-Sierra**, www.pmc-sierra.com.

Gold-covered "microlens" could be breakthrough in infrared imaging

Researchers at Rensselaer Polytechnic University have developed a lensless, gold-covered "microlens" that they believe will lead to breakthroughs in image-signal and infrared-imaging strength. The research uses the properties of nanoscale gold to "squeeze" light into tiny holes in the surface of the device.

The study demonstrates success in enhancing the signal of an infrared detector without also increasing the noise, according to project leader Shawn-Yu Lin, professor of physics at Rensselaer and a member of the university's Future Chips Constellation and Smart Lighting Engineering Research Center.

"We have shown that you can use nanoscopic gold to focus the light entering an infrared detector, which in turn enhances the absorption of photons and enhances the capacity of the embedded quantum dots to convert those photons into electrons," he says.

Researchers establish the detection ability of an infrared photodetector by determining how much signal it receives and dividing that signal by the noise the detector receives. Photodetectors currently employ MCT (mercury-cadmium-telluride) technology, which has a strong signal but long exposure or low-signal imaging. The study creates a plan for developing QDIPs (quantum-dot infrared photodetectors) that can outperform MCTs.

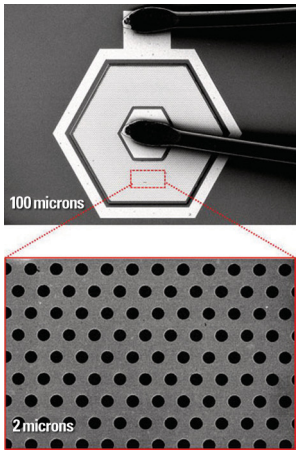

Rensselaer Polytechnic Institute Professor Shawn-Yu Lin has developed a new nanotechnology-based "microlens."

The long, flat surface plasmon QDIPs have countless holes, measuring 1.6 microns in diameter and 1 micron deep, on the surface. Approximately 50 nm of gold covers the solid surface of the structure. Quantum dots—nanoscale crystals with unique optical and semiconductor properties—fill each hole. Properties of the QDIP's gold surface help to focus incoming light directly into the microscale holes and concentrate that light in the pool of quantum dots. That concentration strengthens the interaction between the trapped light and the quantum dots and in turn strengthens the dots' ability to convert those photons into electrons. The end result is that the device creates an electric field as much as 400% stronger than the raw energy that enters the QDIP.

The effect is similar to what would result from covering each tiny hole on the QDIP with a lens but without the extra weight and the hassle and cost of installing and calibrating millions of microscopic lenses.

Lin's team also demonstrated that the nanoscale layer of gold on the QDIP neither adds noise nor affects the device's response time. "Within a few years, we will be able to create a gold-based QDIP device with a 20-fold enhancement in signal from what we have today," Lin says. "It's a reasonable goal and could open a new range of applications, from better night-vision goggles for soldiers to more accurate medical-imaging devices."

Rensselaer published the results of the study online (**Reference 1**).

—by Suzanne Deffree

▷**Rensselaer Polytechnic Institute**, www.rpi.edu.

REFERENCE

1 Chang, Chun-Chieh, Yagya D Sharma, Yong-Sung Kim, Jim A Bur, Rajeev V Shenoi, San-jay Krishna, Danhong Huang, and Shawn-Yu Lin, "A Surface Plasmon Enhanced Infrared Photodetector Based on InAs [Indium-Arsenic] Quantum Dots," *Nano Letters*, April 20, 2010, pg 1704, http://pubs.acs.org/doi/abs/10.1021/nl100081j.

KOZIO HIGHLIGHTS VALIDATION AND TEST SOFTWARE

Kozio recently introduced a software suite for design validation, manufacturing test, and in-field test for embedded systems. The software provides broad datapath coverage and standardized diagnostic tests across functional areas and supports fast, automatic troubleshooting. According to the company, users can save $100,000 in test-development and debugging costs per project and see a three-month reduction in time to market.

The suite helps PCB designers contend with increasingly complex boards. Mentor Graphics figures show, for example, that although average board sizes have decreased over the last 15 years from 101 to 75 in.^{2}, component counts have increased from 649 to 3399, the number of component pins has increased from 4214 to 13,505, and the number of pin-to-pin connections has increased from 5190 to 10,960. Meanwhile, designers are contending with tight schedules, limited resources, design and manufacturing silos, and the need to deal with remote teams.

For design validation, the tools support interactive hardware debugging, fault isolation, characterization, and regression testing. For manufacturing test, they support parallel test and IP (intellectual-property) protection, and contract manufacturers can adapt them to their needs. In the field, the tools support built-in self-test and diagnostics. The tools provide coverage of memory; data buses; user interfaces; displays; cameras; and audio, networking, and wireless functions.

—by Rick Nelson

▷**Kozio**, www.kozio.com

Embedded Answers

by Robert Cravotta – EDN Embedded Master

Q: When does it make sense to use an RTOS or operating system?

A: The answer to when it makes sense to use an RTOS depends on your system's size and complexity, as well as its real-time, reliability, and interface requirements. A key benefit of using an RTOS is that it encompasses thousands of design decisions (for better or worse) and the project does not incur a delay that it would otherwise need to because the design team avoids the productivity drag of having to consider all of the trade-offs for each of those decisions. Contemporary processors and RTOSs offer coupled resources, such as memory protection units and padded cell virtualization support, which may blur the criteria for when it is appropriate to use or avoid using an RTOS. All but the most cost or power constrained projects may be able to afford the runtime overhead associated with at least a small and fast kernel to gain the productivity benefits of using an RTOS.

An RTOS can simplify incorporating commodity features, and it lets you focus your energy on your value add, where your product really differentiates itself from other products.

Feedback from designers indicates that using an RTOS offers a unified hardware abstraction layer and a bunch of inter-process (task) communication APIs that can make application code more portable. If you are building a system that doesn't require user input or output, and doesn't execute much code, you probably don't need an RTOS. If your project has several developers and is using industry standard I/O ports, using an RTOS will be beneficial unless you have a very simple application or you have enough time and money to write the entire product yourself and exhaustively test all the code.

Using an RTOS supports adding the capabilities that your customers demand, by incorporating additional components of the RTOS, such as support for file systems, protocol stacks, task scheduling, Ethernet, USB, wireless, web access, and touch screen GUI. An RTOS can simplify incorporating commodity features, and it lets you focus your energy on your value add, where your product really differentiates itself from other products. An RTOS can simplify migrating to a new processor or hardware, with your software development team not having to rewrite everything for the new hardware.

It might be useful to make a distinction between a real-time "operating system" and a multi-tasking kernel. But at the heart of any RTOS is a multi-tasking kernel whose primary job is to schedule tasks in response to events. A system consisting of synchronous tasks can use a hand-written co-routine scheduler. When a design indicates using multiple asynchronous threads or processes rather than a simpler synchronous operation, an RTOS kernel is recommended. There is a class of not-so-well-known RTOSs that do not include a file-system or support multi-tasking in the usual way - but it delivers a reliable task scheduling service and a minimum of hardware abstraction.

On the other hand, some developers indicate they would not consider using an RTOS if the project exhibits hard real-time or safety integrity requirements and any time overhead is not acceptable. If you need to demonstrate that your system has been developed to a process that conforms to a particular standard, your choices of RTOS are limited. There are a few RTOSs that are designed to be used in military/aerospace/nuclear/automotive applications, but it's more than a rubber-stamping exercise and it may require you to working closely with the RTOS vendor.

For more answers to your Embedded questions, visit
www.embeddedmaster.edn.com.

THE EMBEDDED DEVELOPER

Microsoft

Windows Embedded

BUILT WITH WINDOWS 7 TECHNOLOGIES

Vol. 5

A MALFUNCTION IN THE SYSTEM COULD COST THE PLANT MILLIONS...

WINDOWS® EMBEDDED OFFERS A HIGHLY RELIABLE PLATFORM, WITH THE LEVEL OF PERFORMANCE YOU NEED TO HELP DELIVER CONNECTED DEVICES THAT STAND OUT.

WHICH WINDOWS® EMBEDDED-PLATFORM CAN HELP YOU DELIVER STANDOUT DEVICES?
FIND OUT AT WINDOWSEMBEDDED.COM/DEVICESTORIES

BY HOWARD JOHNSON, PhD

7% solution

In 1983, my mentor, Martin Graham, PhD, had me build a Wheatstone bridge for measuring the common-mode impedance of certain twisted-pair cables, now known generically as Category 3 UTP (unshielded twisted pair). My setup required some closely matched carbon-composition resistors with values accurate to within ±½%.

The measurements would be taken at fairly high frequencies covering the 10- to 100-MHz range. At such frequencies, even the parasitic series inductances and shunt capacitances of the resistors must match. That requirement rules out the use of trimming potentiometers to meet the stringent accuracy requirement.

I needed a few bull's eye, hit-the-spot, on-the-money, perfect resistors. Looking at the lab stock available on that day, I found no high-precision carbon-composition resistors. There were some 2% metal-film resistors in stock, but I knew that the manufacturers of those parts sometimes etched serpentine patterns in the metal film to elongate the resistor, thus increasing the parasitic series inductance to levels unacceptable for my application. Carbon-composition resistors are made in a simple cylindrical shape that is ideal for high-frequency use.

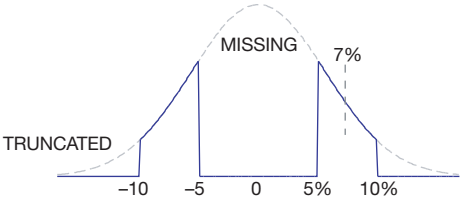

Figure 1 This histogram of 10%-resistor values reveals how 5% resistors are "made."

Imagine my surprise when, after an hour of labor and after checking 300 resistors, I found that none fell within my 1% selection window.

The only carbon-composition resistors I found had a 10% tolerance. I decided that they might work if I hand-selected values good enough for my purpose. I reasoned that out of 100 parts rated at 10% tolerance, about 10 should fall within 1% of the advertised value and that even more would do so if the distribution were centrally clumped. From those 10 parts, I hoped to select a couple of pairs suitable for my setup.

Imagine my surprise when, after an hour of labor and after checking 300 resistors, I found that none—absolutely none—fell within my 1% initial selection window.

Mathematically, if the component values were distributed evenly across the whole tolerance band of ±10%, the probability of any one resistor's falling within a ±1% selection window should be one out of 10, or 0.1. The probability of any one resistor's falling outside the selection window then equals the complement of that value, 0.9. If you repeat the experiment 300 times, the probability of all the resistors' falling outside the selection window equals $(0.9)^{300}=1.8\times10^{-14}$. It seemed to me inconceivable that such a low-probability event could ever actually occur in my sample of 300 parts. Monkeys striking random keys on a typewriter could more likely compose a sonnet in the mean time between experimental failures of that magnitude.

Perplexed, I sought guidance from Martin. I found him in the company cafeteria enjoying a meatloaf sandwich. On a ketchup-stained napkin, he patiently drew an odd-looking curve (**Figure 1**).

The drawing complete, Martin said, "A 10% carbon-composition resistor is made in a somewhat slipshod manner. The manufacturer tries to get it right, but some of the variables are just too difficult to control. They make up a batch, test them all, and then throw away the bad ones. What's left is a distribution of values truncated on either side at the ±10% limits. The other main feature of the distribution is the big gap-toothed section in the middle. That's where they pulled out all the good parts and sold them at a higher price with a ±5% tolerance. How else do you think they make 5% resistors?"

My jaw hit the floor when I grasped how perfectly his explanation matched my results. He paused and then passed along another point of wisdom: "Design your circuit to use values that lie 7% away from the nominal standard values, either higher or lower, and you'll find plenty of those in the bins when you do your hand selection." EDN

Howard Johnson, PhD, of Signal Consulting, frequently conducts technical workshops for digital engineers at Oxford University and other sites worldwide. Visit his Web site at www.sigcon.com, or e-mail him at howie03@sigcon.com.

Honeywell & TTI – a powerful new team working for you!

Honeywell and TTI together are proud to bring you the industry leading line of sensors and switches, designed to ensure the integrity of your product.

With quality products from Honeywell and reliable, assured delivery service from TTI, you can be confident you have a total solution to meet your switch and sensor needs.

Call TTI today for your free Honeywell Solutions Guide

Honeywell

tti®

NEVER SHORT ON SOLUTIONS

1-800-CALL-TTI

www.ttiinc.com

NEDA Supplier Authorized Distributor

BY MIKE SALAS • SILICON LABS

LOW POWER: a chip and system-design primer

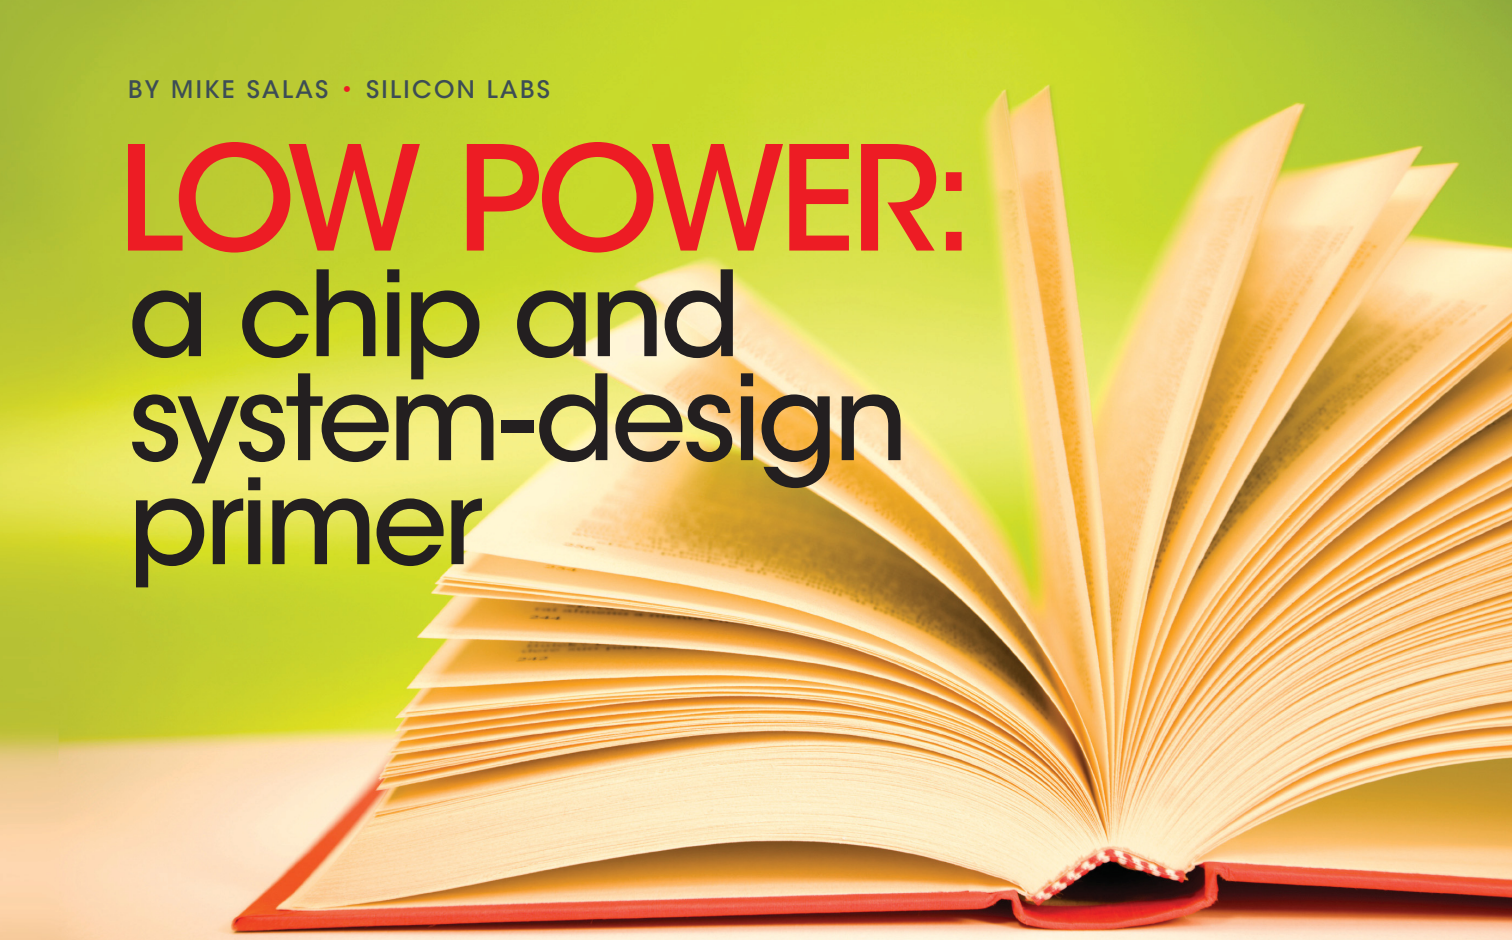

PARSING THROUGH THE CONFLICTING POWER-CONSUMPTION CLAIMS OF VARIOUS MICROCONTROLLER VENDORS CAN BE CHALLENGING. FOR MOST APPLICATIONS, A DUTY-CYCLE-DETERMINED FUNDAMENTAL POWER-CONSUMPTION EQUATION CAN EASILY CUT THROUGH THE CLUTTER.

Reducing power consumption has a major impact on every aspect of our lives. At a macro level, the benefits have been well-documented: lower electric bills for consumers, reduced load on utilities, and fewer batteries in landfills. In short, saving power is good for both the environment and the pocketbook. Due to the growing use of electronics worldwide, reducing power consumption must begin at the microchip level. Power-saving techniques that engineers have designed in at the chip level have a far-reaching impact, especially when involving microcontrollers that serve as the engines behind most of these electronic devices.

From a system-design perspective, identifying which microcontrollers are truly low-power requires designers to navigate through the myriad claims of various semiconductor vendors. Because of the varying and confusing metrics vendors use, this objective is a complicated task. This article briefly describes the main factors that you need to consider when analyzing competitive microcontroller alternatives. At a basic level, you can define microcontroller power consumption as the sum of active-mode power and standby, or sleep-mode, power. However, another important metric to keep in mind is the amount of time it takes for a microcontroller to move from a standby state to an active state. Because the microcontroller cannot do any useful processing until all of its digital and analog components are fully settled and operational, adding in this wasted power is important when calculating total power consumption. Thus, the device's total power consumption is the sum of its active-mode power, standby power, and wake-up power.

Because every application is different, system designers have a tendency to give more weight to some of these elements than others. For example, some applications, such as water meters, spend most of their time in a standby state, so their long duty cycles require low standby-power consumption. Other applications, such as data loggers, often enter and exit active states, so it is critical to limit the time they spend in wake-up-transition modes (**Figure 1**). However, a vendor developing a compelling microcontroller cannot attempt to guess which of these metrics is most

336 Volts of Green Engineering

MEASURE IT – FIX IT

Developing a commercially viable fuel cell vehicle has been a significant challenge because of the considerable expense of designing and testing each new concept. With NI LabVIEW graphical programming and NI CompactRIO hardware, Ford quickly prototyped fuel cell control unit iterations, resulting in the world's first fuel cell plug-in hybrid.

MEASURE IT

Acquire
Acquire and measure data from any sensor or signal

Analyze
Analyze and extract information with signal processing

Present
Present data with HMIs, Web interfaces, and reports

FIX IT

Design
Design optimized control algorithms and systems

Prototype
Prototype designs on ready-to-run hardware

Deploy
Deploy to the hardware platform you choose

Ford is just one of many customers using the NI graphical system design platform to improve the world around them. Engineers and scientists in virtually every industry are creating new ways to measure and fix industrial machines and processes so they can do their jobs better and more efficiently. And, along the way, they are creating innovative solutions to address some of today's most pressing environmental issues.

>> Download the Ford technical case study at **ni.com/336**

800 258 7018

NATIONAL INSTRUMENTS™

important but instead designs a system from the ground up that focuses on minimizing any power consumption. Accomplishing this objective requires strong mixed-signal expertise to address both the architectural-level and the circuit-level challenges necessary to minimize power in both the analog and the digital domains.

ACTIVE-MODE CURRENT

For a CMOS logic gate, the following **equation** yields dynamic power consumption: $C \times V^2 \times f$, where C is the load capacitance, V is the supply voltage, and f is the switching frequency. The capacitance term is a function of the design and processing technology, and the frequency term is a function of the application's processing requirements. However, as you can see from the **equation**, the supply voltage has a disproportionate impact on the overall power that the microcontroller consumes. Therefore, adding voltage regulation to the microcontroller design can yield significant active-mode-power savings by providing a lower, steady supply voltage to the microcontroller's circuitry. Switching-type converters may be a possible approach for this application, but they better suit regulator environments requiring large voltage-conversion ratios. However, in battery-type applications in which the average voltage-conversion ratio approaches 1-to-1 at the end of the battery's life, a better approach would be to add an on-chip low-dropout linear voltage regulator because it offers acceptable efficiency with less complexity and lower cost than a switching approach.

To illustrate the benefits of using a low-dropout regulator, it is helpful to restate the CMOS dynamic power equation as $C \times V^2 \times f = V \times (C \times V \times f) = V \times I$, where the dynamic current, I, equals $C \times V \times f$. It is common to normalize the dynamic current to a frequency of 1 MHz and a particular supply voltage. For example, one recently introduced ultra-low-power microcontroller has a dynamic current consumption of 160 mA/MHz at 1.8V. Without supply regulation, this metric would increase to $160 \times (3.2/1.8) = 284$ mA/MHz when the supply voltage is 3.2V. With a low-dropout regulator, the battery current remains at 160 mA/MHz across the entire supply range (**Figure 2**).

AT A GLANCE

- Reducing power consumption must begin at the microchip level.
- Identifying which microcontrollers are truly low-power requires designers to navigate through the myriad claims of various semiconductor vendors.
- Adding voltage regulation to the microcontroller design can yield significant active-mode-power savings by providing a lower, steady supply voltage to the microcontroller's circuitry.
- An on-chip low-dropout linear voltage regulator offers acceptable efficiency with less complexity and lower cost than a switching approach.
- An innovative way to reduce power and environmental impact is to convert the design to a single-battery configuration in which the battery operates at 0.9V to the end of its useful life.

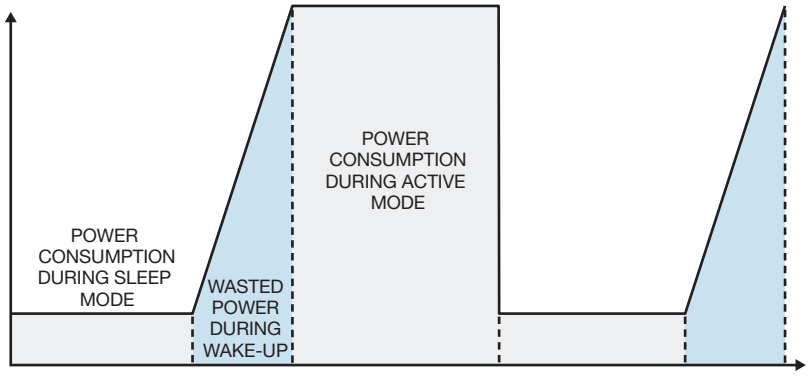

Figure 1 Curtailing wasteful wake-up latencies is critical to minimizing overall power consumption for applications that frequently move into and out of active operating modes.

You can use this advanced power architecture to maintain a constant active current over the full operating-voltage range, and it can help you achieve a significant savings in power consumption. Therefore, it is important to determine the microcontroller's current consumption when operating across the entire operating-voltage range, not just at the 1.8V minimum operating condition that microcontroller vendors commonly quote. Quoting an optimistic current number that assumes anything less than a typical voltage supply does not accurately reflect how applications find use in the real world.

As an example, in systems requiring two AA or AAA coin-cell batteries, the batteries most often operate near their initial 3V voltages. Therefore, the quoted 1.8V specification can be deceiving because, from this perspective, most microcontrollers consume approximately 50% more power than the amount that their vendors commonly quote. Furthermore, because power consumption is directly proportional to switching frequency, system designers should normalize the quoted current numbers down to a current-per-megahertz basis. By combining these two factors, you can perform a side-by-side comparison of microcontrollers using current consumption per megahertz at 3V.

Some vendors attempt to confuse the issue by equating megahertz to system-clock speed when the value that is truly meaningful is instruction-clock speed. This substitution is deceiving because system clocks can operate at least twice as fast as they can execute instructions, thereby at least doubling their effective power consumption. It is therefore important to normalize specifications to instruction-clock speed. By doing so and by using a typical supply voltage, you can properly derive the actual active-mode current-consumption budget.

STANDBY CURRENT

Achieving maximum energy efficiency and, therefore, battery life requires ensuring that each microcontroller task consumes the minimum possible current at the minimum possible voltage for the shortest possible duration, so that the device spends most of its time in a low-power sleep mode. In

Renewable Energy Applications: Reliability and Safety with Fiber Optic and Optocoupler Isolation Technology

Introduction

Climate change, rising carbon dioxide levels, environmental pollution and volatile fossil fuel prices have driven electric power generation producers, governments, and individual consumers to adopt alternative energy generation methods such as solar and wind. Governmental renewable energy support has been incorporate into strategic energy plans worldwide. During 2009 at least 73 countries had energy policy targets and more than 64 countries had policies to promote renewable power generation. Renewable energy is no longer an area dominated by research, but rather a large commercial market with high adoption and installation growth rates.

Renewable energy growth in 2008 was substantial and 2009 activity points toward continued progress. The United States ended 2008 with 25 GW of wind power, eclipsing Germany's 24 GW capacity. China's total wind power doubled for the fifth year in a row, ending the year above 12 GW, exceeding China's 2010 development target of 10 GW two years ahead of plan. More than 80 countries around the world now have commercial wind power installations.

With most 2008 offshore wind generation capacity in Europe, capacity reached nearly 1.5 GW. The United Kingdom became the offshore wind power leader in 2008 but other countries, such as the United States, have plans for new offshore facilities.

Grid-connected solar photovoltaic plants are the fastest growing power generation technology, with a 70-percent increase in existing capacity to 13 GW in 2008. This was a six fold increase in global capacity since 2004. Including off-grid applications, total 2008 photovoltaic generation increased to over 16 GW. In 2008, utility-scale solar PV power plants—plants over 200 kW—were estimated to have grown by 80%, to 1,800, since 2007. The added plants totaled over 3 GW, a tripling of existing 2007 capacity.

Utility photovoltaic plant additions were a major 2008 focus in Spain but facilities also went into the Czech Republic, France, Germany, Italy, Korea, and Portugal. New solar photovoltaic plants are planned and under development throughout Europe as well as China, India, Japan, and the United States.

Given the high growth rates and importance of renewable energy, Avago focused its fiber optic and isolation product portfolios and product development efforts toward renewable energy applications. Digital optocouplers, gate drivers, isolation amplifiers and solid state relays provide safety isolation and insulation along with their basic function. Fiber optic components, besides providing isolation, provide EMI resistant, long distance communication channels in wind turbine farms and panel-to-panel communications in solar photovoltaic systems. These control and communication channels are critical for safety and optimal power efficiency in distributed solar systems and especially for offshore wind farm installations.

Avago Renewable Energy Solutions

Avago Technologies offers fiber optic transmitters, receivers and transceivers, as well as IGBT and Power MOSFET gate drivers and optocoupler isolation products for wind turbine, wind farm and solar electric power generation.

Avago fiber optic components, isolated gate drivers, isolation amplifiers and digital optocouplers are deployed in wind turbines, wind farms and solar photovoltaic energy farms today.

Applications include:

- Power rectifiers and inverter power assemblies
- Control and communication boards
- Turbine control unit condition monitoring systems
- Wind farm networks
- Sun tracking control
- Communications boards
- Solar farm substation automation and protection relays
- Single- and three-phase inverters
- Battery chargers
- Motor control
- DC-to-DC converters,
- System control and monitoring

Solar Energy: Fiber Optic Solutions

Avago fiber optic transceivers provide control and communication links within a solar energy system and within the solar farm. Resistant to EMI and able to withstand rugged environments, Avago transceivers support POF, HCS and Multimode fiber. Plastic cable supports transmission up to 50 m. For longer distances 200 µm HCS and 62.5/125 µm fiber will support up to 2,700 m links. Maximum transmission distances, depending on data rates, are up to 80 km.

Data rates from 1/5/10/20/160 MBd are available with many different connector styles: LC, SC, SFP and MT-RJ.

Figure 1. Fiber data links in solar energy farm

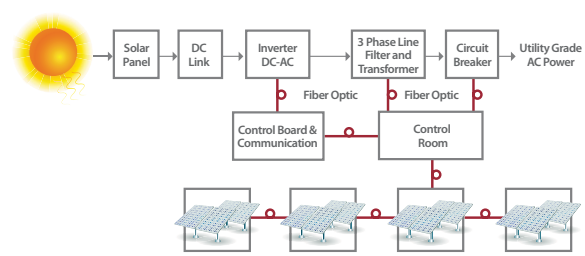

Wind Turbine Networking, Control and Communications: Fiber Optic Solutions

In wind turbine applications access for maintenance and repair is more limited than in land-based solar farms. This is especially true for off shore wind generation installations. In addition, large mechanical loads and changing local weather conditions must be monitored in near real time for optimal power generation, safety and control. Long term reliability and quality are very important component selection factors.

Inside the turbine nacelle, short range fiber optic connections connect power generation subsystems and the turbine pitch controller to the main system controller. Avago offers a broad portfolio of fiber-optics transmitter and receiver like the

HFBR-0500Z family which are used in Wind-Turbines since many years. For Wind-Turbines >2 MW the main choice inside the turbine is fiber optics because of EMI, lightning and galvanic isolation. Longer range fiber links connect to tower base and the wind farm network.

In turbine monitoring and solar systems, Fast Ethernet transceivers with DMI and SC-RJ connectors, such as the Avago AFBR-5978Z, are popular as they operate at 650 µm over POF and HCS cable and feature a –25°C to 85°C temperature range. A demonstration kit for the AFBR-5978Z (AFBR-0978Z) is available from any Avago sales offices or representative.

Our new multimode, small form factor (SFP) HFBR-57E5APZ Fast Ethernet transceiver with DMI requires only 611 mW and operates from 3.3 V.

Figure 2. Fiber optic communication links in wind turbines

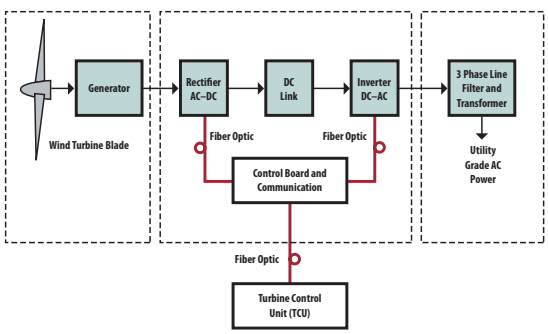

Solar and Wind Power: Application Specific Optocouplers, Gate Drivers and Isolation Amplifiers

As the leading optocoupler supplier, Avago develops application specific isolated power MOSFET gate and IGBT drivers for power applications. Featuring high peak output current, high CMR, wide temperature range, low supply current, surface mount packages and integrated protection circuits they are ideal for motor control, inverter and battery charging circuits.

The ACNW3190 isolated gate driver has a 5 A peak drive current and under voltage lockout (UVLO) protection. Other devices have 0.4 A to 2.5 A peak drive current. For maximum reliability and design flexibility, Avago devices have protection features such as integrated VCE detection, UVLO, "soft" IGBT turn-off, isolated open-collector fault feedback and active Miller clamping for maximum design flexibility and circuit protection. A Miller clamp controls the Miller capacitance current during high dV/dt transitions and can eliminate the need for a negative supply voltage.

Figure 3. Isolation and insulation solutions in wind-solar power generation systems include inverters, charge controllers and communications.

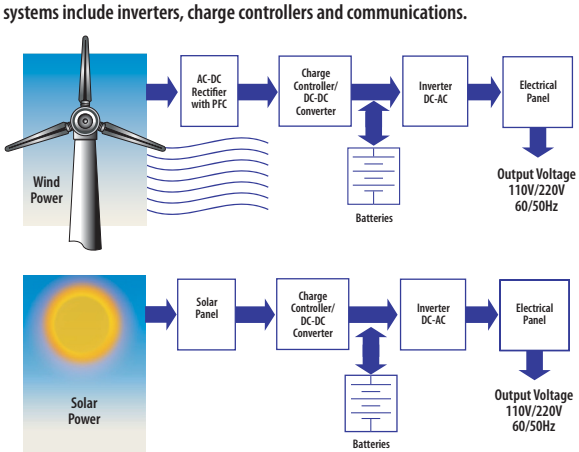

With the Avago analog isolation amplifiers, current and voltage monitoring in motor control, inverter and battery applications is simple, safe and accurate. Phase currents and DC link voltages are easily measured while maintaining isolation. The amplifiers have working voltages of up to 1140 V for safety. Gain accuracy ranges from ±1% to ±5%.

The ACPL-796J 1-bit, second-order sigma-delta modulator converts an analog input signal into a high-speed data stream with galvanic isolation based on Avago's optical coupling technology. The ACPL-796J operates from a 5 V power supply with dynamic range of 80 dB with an appropriate digital filter. The differential inputs of ±200 mV, full scale ±320 mV, are ideal for direct connection to shunt resistors or other low-level signal sources in motor phase current measurement applications.

Solar and Wind Power: Digital Optocouplers

Wind control systems must accommodate varying wind conditions and avoid run-to-failure by monitoring subsystems accurately. High voltage isolation, current isolation and immunity to EMI are critical for reliable and robust control systems. Solar and wind farms face harsh environmental conditions and operate over large physical areas susceptible to lightning and extreme temperatures. Avago digital optocouplers have the insulation and isolation specifications to make command and control systems safe and reliable. All our digital optocouplers are qualified to all international safety organization insulation and isolation specifications. Many devices have working voltages that exceed 1768 Vrms and isolation voltage up to 7500 Vrms are offered.

Operating at 10 MBd, multi-channel and bi-directional optocouplers (ACSL-6210/6310/64x0 in Dual/Triple/Quad channel) maximize design options and conserve PCB area. For extreme environments, the ACPL-M61U features –40°C to 125°C operation. The ACPL-M61L/061L/064L/W61L/K64L are ultra low power optocouplers suitable for energy efficient systems. Higher speed devices such as the ACPL-772L/072L operate at 25 MBd and from –40°C to 105°C.

Our photo MOSFETS, or solid state relays (SSRs), operate with voltages from 60 V to 600 V and with currents up to 2 A for energy storage management systems.

Figure 4. Isolated gate drive, digital optocoupler, level shifter and isolation amplifier applications.

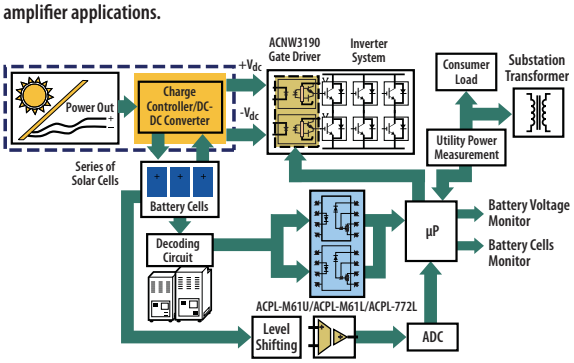

Summary

Innovative fiber optic and optical isolation solutions have been available from Avago for over three decades. As a major contributor to the work of international standards organizations, device reliability, quality and safety are always part of any Avago design. We offer the industry's best isolation technology, specifically designed and manufactured to meet the stringent requirements of applications in power generation systems.

Around the world our customers value what we do. We provide the components and technology to make their visions for the renewable energy market expand and prosper.

Contact us for your design needs at: **www.avagoresponsecenter.com/401**

AVAGO TECHNOLOGIES

some applications, the sleep-mode current is the parameter most responsible for overall energy consumption. However, engineers often overlook the fact that leakage current is the primary limitation on the absolute minimum sleep current a microcontroller can achieve. For example, a 20-input device that has an input-leakage-current specification of 100 nA could consume as much as 2 µA of power during sleep mode.

A number of factors, the most important of which is process technology, affect leakage current. In some cases, vendors use a 0.25- or 0.35-micron process technology to reduce the leakage-induced sleep current, but this choice comes at the expense of a higher active current. In other cases, microcontroller vendors use 0.18-micron or smaller process technologies to reduce active-mode current, but this choice comes at the expense of higher leakage currents. One approach to this problem is to apply mixed-signal expertise to implement an advanced power-management unit that limits leakage and enables ultralow sleep current regardless of the underlying process technology in use. When using process technologies of 0.25 micron or smaller, minimizing sleep-mode current requires cutting power to the digital core. Modules that operate in sleep mode, such as power-management circuits, I/O-pad cells, and real-time clocks, must operate from the unregulated voltage supply to avoid burning additional current in a low-dropout regulator. Cutting power to the digital-core logic also prevents its off-state leakage from contributing to the sleep-mode current; however, the microcontroller must preserve RAM contents and the state of all registers during sleep mode so that execution can resume where it left off.

You can perform this preservation using a low-current sleep-mode latch-bias scheme or retention latches that can hold the state in sleep mode without significant leakage. The microcontroller also requires some form of continuous supply-voltage monitoring, such as brownout detection, to reset the device in the event that the supply voltage drops below the minimum retention voltage, which could corrupt the operating state. It is therefore important to examine the underlying leakage-current specifications to determine which microcontroller vendors have applied their mixed-signal expertise toward solving this complex problem.

Designers should also consider the fact that most vendors offer many standby-current options. Most suppliers highlight their absolute lowest sleep-mode current, which often corresponds to the current the system consumes when the real-time clock and brownout detector are disabled. Some vendors go a step further and quote a shutdown-mode current that does not retain memory and requires a reset to wake up, which in general is not a practical mode. Therefore, because most applications require full RAM and register retention, it is important to perform side-by-side comparisons using the standby/sleep-mode current with real-time clock and brownout disabled with RAM retention, standby/sleep-mode current with real-time clock disabled and brownout enabled, and standby/sleep-mode current with real-time clock and brownout enabled. You can then use the correct values when calculating the overall standby-mode power budget based on the duty cycle of the application.

IMPLEMENT AN ADVANCED POWER-MANAGEMENT UNIT THAT LIMITS LEAKAGE AND ENABLES ULTRALOW SLEEP CURRENT.

WAKE-UP ENERGY

Systems that use sleep modes can waste a significant amount of power waking up the microcontroller and preparing it to acquire or process data. In some applications, a microcontroller can often use just as much energy when coming out of standby as when the device is fully processing data. Therefore, it is important to design a microcontroller to wake up and settle in a short time to minimize the amount of time it spends in an energy-wasting state. The microcontroller should be able to exit sleep mode from either an external

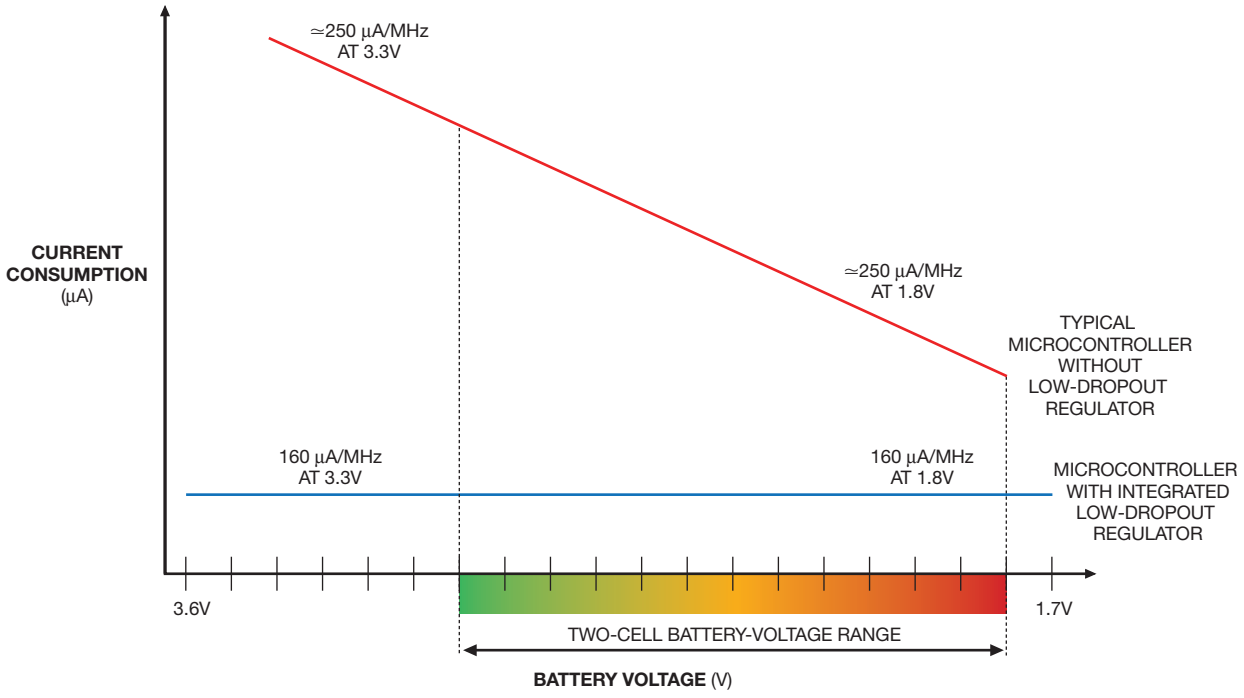

Figure 2 A low-dropout regulator can hold microcontroller power consumption constant across a wide supply-voltage range.

trigger event or an internal timer. The most flexible periodic-wake-up source is a real-time clock that can operate from an external crystal oscillator for applications requiring accurate timing or from a low-frequency internal oscillator that eliminates the need for a crystal in applications not requiring high accuracy. Avoid using a slow-starting crys-

tal oscillator for the high-speed system clock; an accurate, quick-starting, on-chip oscillator is a better alternative.

Because many products periodically wake up to sample an input using on-chip ADCs, it is also important to allow enough time for both the digital circuitry to wake up and the analog circuitry to settle before beginning to make valid measurements. The start-up behavior of the analog modules can have a major impact on the amount of time the system spends in active mode; voltage regulators or references using external decoupling capacitors can take many milliseconds to settle. Microcontroller vendors sometimes quote only the wake-up times for the digital circuitry and ignore the time it takes for the analog circuitry to settle. Therefore, it is important to analyze the overall wake-up and settling time for both the digital and the analog circuitry to factor in the true cost of this wasted energy.

OTHER CONSIDERATIONS

Other ways also exist to further reduce power in a system. For example, you can use two AA or AAA batteries in some configurations because microcontrollers often can typically operate with voltages as low as 1.8V with fewer functions. For example, they may lack an ADC or have a slower instruction clock. An innovative way to reduce power and environmental impact is to convert the design to a single-battery configuration in which the battery operates at 0.9V to the end of its useful life. To enable this approach, a microcontroller must integrate a highly optimized dc/dc converter that can operate to the lowest usable voltage of the battery, which, in the case of alkaline chemistry, is 0.9V. This approach can also save the supplier, the consumer, or both the cost of a battery.

Another method of reducing power is to use highly integrated microcontrollers that include ADCs, DACs, and other peripherals because the microcontroller can have control over enabling and disabling these peripherals when the applications require them. For example, some microcontrollers offer specialized low-power ADCs with burst modes that can take analog measurements when the CPU is off in an effort to further minimize power consumption. EDN

AUTHOR'S BIOGRAPHY

Mike Salas joined Silicon Laboratories in 2008, serving as business-development director. In his current role as microcontroller-marketing director, he oversees the company's global microcontroller business. Before joining Silicon Labs, he founded and served as chief executive officer of Layer N Networks, which later became Britestream Networks, a private, venture-capital-backed company that offered semiconductor products for networked security. From 1994 to 2000, he served as business-development director at PMC-Sierra, where he helped oversee the company's growth from early start-up to a dominant networking-semiconductor specialist. Previously, Salas held various management and technology positions at Texas Instruments. Salas obtained a bachelor's degree in electrical engineering from the University of Texas—Austin.

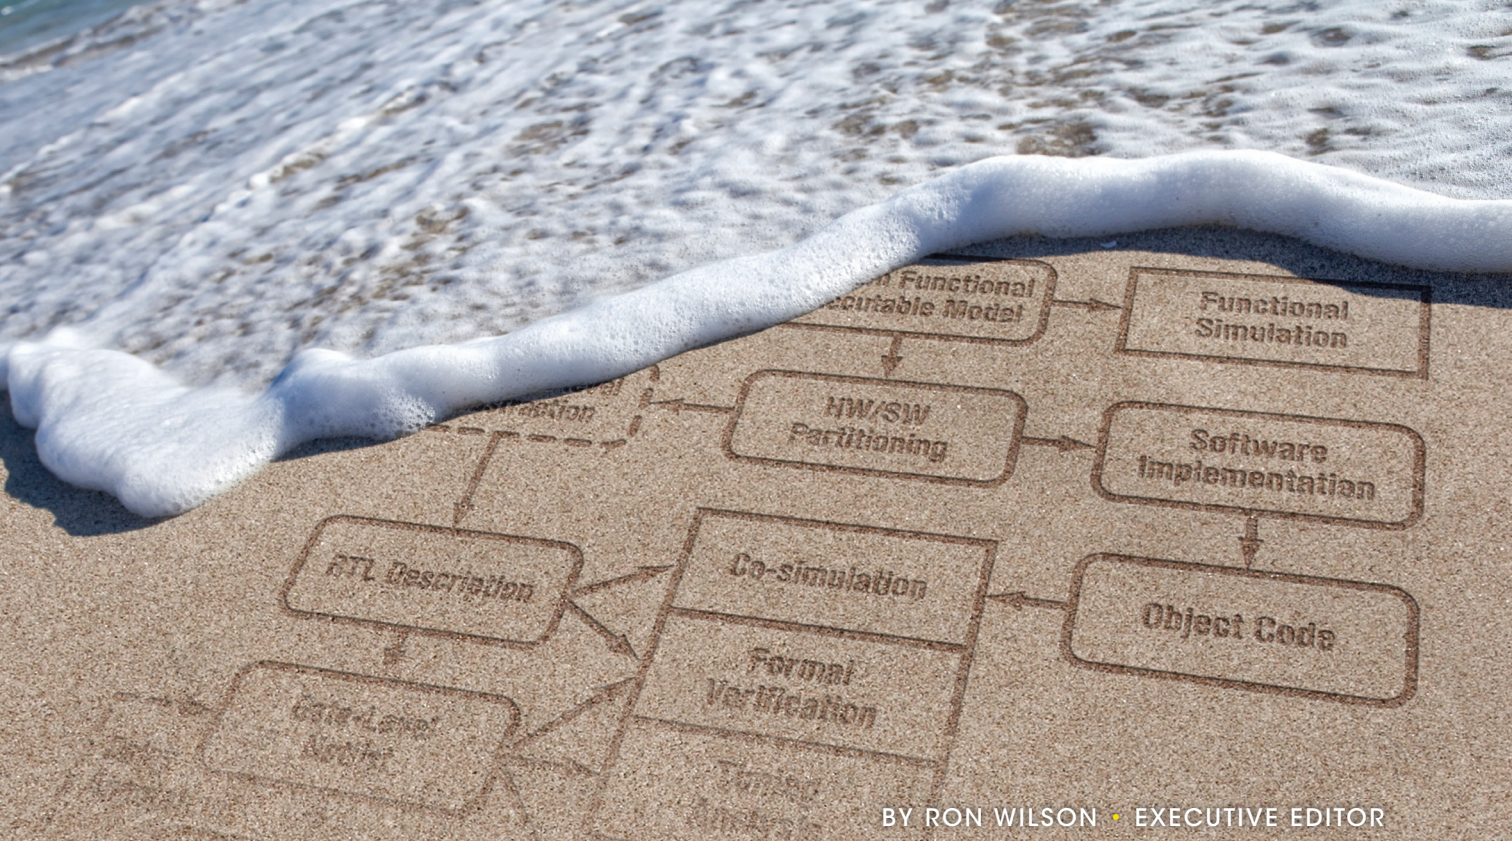

BY RON WILSON • EXECUTIVE EDITOR

ALTERING THE SOC-DESIGN FLOW

POWERFUL FORCES ARE AT WORK, RESCULPTING SOC-DESIGN METHODS INTO NEW SHAPES.

SOC (SYSTEM-ON-CHIP)-DESIGN TEAMS at the leading edge of their markets say that "business as usual" is no longer the case. Powerful technical and business forces, seemingly independent of the EDA vendors' road maps, are resculpting SOC-design methods into new shapes, profoundly different from the best practices of only a few years ago. The change is painful for many architects, designers, and managers. To cling to the past is to embrace failure, however.

The driving forces of this change are several. Financial and geographic realities have mandated a growing dependence upon third-party IP (intellectual property) and have attenuated the feedback loop from downstream issues to RTL (register-transfer-level) corrections. Complexity—especially in power and clock networks of aggressively power-managed designs—has forced tasks that formerly resided downstream to cluster early in the design flow. Also, the challenges of advanced processes have influenced both front- and back-end practices.

THE DRIVING FORCES

IP was supposed to be part of the solution, not part of the problem. IP reuse—of everything from I/O controllers to CPUs—has made it possible to both disperse and shrink design teams. The pervasive use of IP has changed the nature of the design flow, however. The flow formerly comprised building behavioral requirements, reducing them to RTL, synthesizing a netlist, and implementing it in cells. It has now become a process of assembling and imposing closure on an ad hoc set of complex, increasingly immutable, and opaque functional blocks. When designers encounter problems with integration or closure, the original IP developer is often the only one who can help.

Although IP reuse has helped with the sheer size of designs, it has not helped with other facets of complexity. This situation is especially true for power management. Clock gating is a mandatory design step for reducing dynamic power. But it has made SOCs' clock networks so complex that the clock trees are, in essence, other signal networks, requiring extraction, timing, power, and signal-integrity closure. And the use of voltage islands, power gating, and DVFS (dynamic voltage/frequency scaling)—just coming into use in most design teams—promises to similarly complicate power grids.

Finally, the processes themselves are forcing change. Despite the heroic efforts of process engineers and cell-library developers, the complexities of advanced processes have, by the 65-nm node, begun to tunnel through the custom/cell-based barrier to present themselves to chip designers. "Our memory-compiler designers have had to deal with process variations, with the weak drive strengths of cells, and with increasingly complex DFM [design-for-manufacturing] rules," says Lisa Minwell, director of technical marketing at Virage Logic. All these issues are now facing chip designers in cell-based flows. These forces have combined not merely to make design more difficult, but also to change the underlying approach.

THE CRUCIAL BEGINNING

Open-Silicon recently executed a 100 million-gate wireless-networking SOC. The company executed this design in TSMC's (Taiwan Semiconductor Manufacturing Co's) 65-nm CMOS process. "The key to this design was upfront planning," says Taher Madraswala, Open-Silicon's vice president of engineering. Open-Silicon partnered on the physical design of the chip with ASIC-design-services company Brite Semiconductor, using requirements and RTL from HiSilicon, a fabless semiconductor vendor. "This was pretty much a top-down design," Madraswala says, noting that clock layout significantly drove the upfront work.

AT A GLANCE

- IP (intellectual-property) reuse, power management, and advanced processes are altering the SOC (system-on-chip)-design flow.
- More work is shifting to the planning stage of the design.
- Power management can have a disproportionate impact on verification and physical design.
- Until tools can catch up, plan on the need for many iterations to get through design closure.

Open-Silicon's task began with understanding the design and performing risk assessment. "This was going to be a very large die with some very long tracks," he says. "So we spent three days in meetings just understanding the clock structure, for instance." Understanding the sources, consumers, and gating structures of the various clocks was a necessary preliminary to block placement. If the team got it wrong, there would be little chance of closing clock timing.

Open-Silicon had to work with multiple instances of IP cores that had essentially fixed pin placement, placing another constraint on block placement. "The problem was repeatability," Madraswala explains. "If you change the orientation of the core, critical routing lengths change, and you get different timing." The team thus performed a preliminary routing of top-level signals, clocks, and I/Os and then used that routing as the basis of the design's partitioning and the placement of the resulting blocks.

"It is very difficult now to partition a design at the system level," says Venkat Mattela, Redpine Signals' chairman and chief executive officer, noting that engineers must do power planning early. In Redpine's design, an extremely low-power 802.11n transceiver for embedded-system applications, the definitions of the modules at RTL made them self-contained entities with respect to the chip's power strategy. The partitioning into modules followed not only functional boundaries but also the boundaries between voltage islands and clock domains. Consequently, the design team could capture the power intent for each block of RTL in a UPF (Universal Power Format) file at the outset of the design.

Other issues can also demand attention early in the design. For example, Vitesse Semiconductor recently developed a 24-port switch SOC with integrated copper PHY (physical-layer) blocks (**Figure 1**). According to Mandeep Chadra, director of design at the company, the power consumption of the PHY blocks dominated the task of estimating how much the designers could integrate because these blocks consume much of the total power. "Throughout the planning process, power remained a major issue, especially since we were targeting a wire-bond package," he says. Far from being an afterthought, that package intruded into the early decisions in chip planning. Without a flip-chip signal-redistribution layer, the arrangement of I/Os on the die had to reflect the pinout of the chip. At these frequencies, the chip had to reflect the layout of the boards that would use it. As a result, the physical layout of a switch box directly influenced the company's floorplan, Chadra adds.

Two issues—power-management strategy and top-level signal, clock, and power routing—emerge in the early stages of chip planning, and EDA vendors have reacted to these changes. Power-aware flows from all the major companies now encourage designers to capture power intent early in standard CPF (Common Power Format) or UPF files, which then guide implementation of power management through synthesis, placement and routing, and verification.

Vendors are paying increasing attention to the need for design teams to have preliminary routing information as early as the partitioning and floorplanning stages of the design. "In the early stages of design, the biggest surprise is congestion," says Pravin Madhani, division manager of Mentor Graphics' place-and-route division. "So people are running their place-and-route tools very

early to check for potential congestion problems." This trend, in turn, is leading place-and-route-tool vendors to extend their tools for use in the preliminary stages of the design.

Surprise congestion issues yield costly consequences. "We encountered a couple of blocks with congestion issues," says Open-Silicon's Madraswala. "We had to go back and get a rewrite on the RTL to deal with them." That step involved another pass through verification, setup, and synthesis for those blocks. Open-Silicon had from the outset, however, cultivated a fast feedback path to the RTL designers at HiSilicon by putting a six-person design team in HiSilicon's location in China.

Congestion surprises with third-party IP can be worse. For instance, what if the IP vendor lacks the resources to make RTL changes on your schedule or if the congestion is at the pins of a hard-IP block? In the worst-case scenario, the SOC team may have to replace the IP vendor. Having the design partitioned and placed consistently with the power strategy and having an early view of top-level routing have thus become mission-critical issues.

SYNTHESIS AND VERIFICATION

The design teams at Open-Silicon, Vitesse, and Redpine don't find synthesis a major problem. Rather, they focus on how to avoid repeated trips through synthesis. "We treated each block of RTL as if it were an independent die," Madraswala says. "Then we focused on exiting each step in the flow for each block at a high-enough level of quality of results. Maybe as a consequence, after clock insertion, we made only one pass through synthesis." Open-Silicon used its synthesis tool to automatically insert clock gating. Otherwise, configuration at the architectural level handles the power management in the chip, according to Madraswala. "There are power islands, but, because the power management comes explicitly through the RTL, we

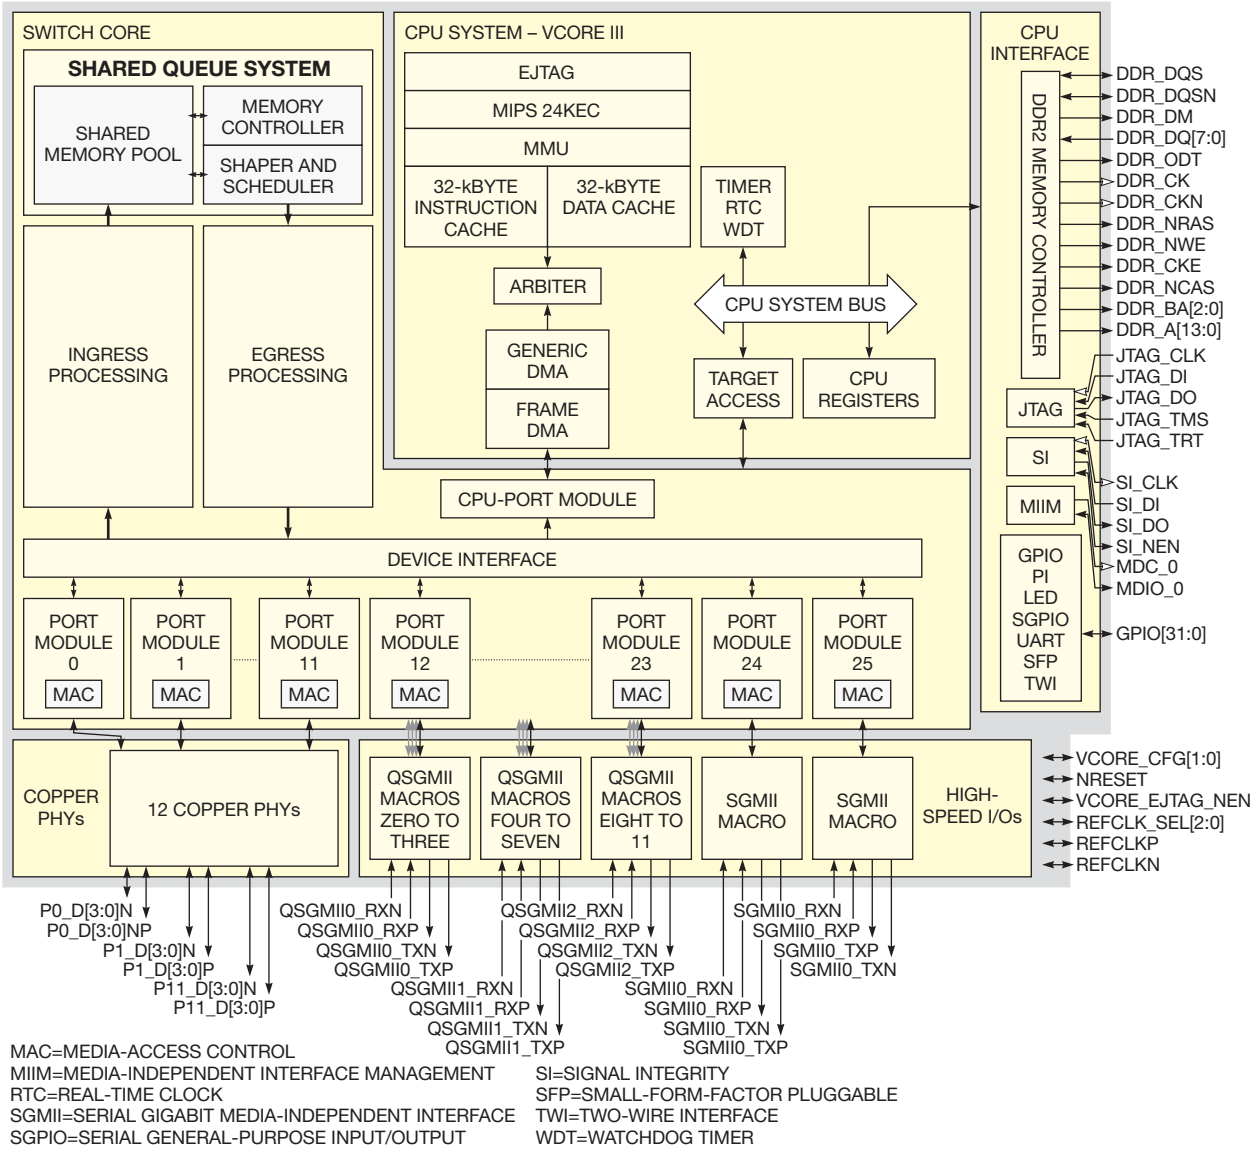

Figure 1 The Vitesse 7427 combines a 24-port switch and MIPS processor with integrated PHYs and MACs.

didn't need anything like CPF," he says. Similarly, the Vitesse design used extensive clock gating, but had only one power-gated block, and Chadra reports no issues with the normal synthesis flow.

Redpine, however, used a more aggressive power-management strategy and pushed the tools harder. That approach had an impact on the design flow (**Figure 2**). Mattela says that, in principle, if you have correctly organized the RTL and accurately captured your power intent, you should be able to feed the RTL, the UPF, and a power-aware library into synthesis and receive a netlist with all the isolators, level shifters, and controls in place. In reality, though, "you push the button, and it doesn't happen," he laments. Everything may look right structurally, but a detailed manual verification with voltage-aware tools may tell a different story.

Verification appears to have adapted to the new order more than has synthesis. With growing complexity, functional verification is starting earlier and at a more abstract level. "We followed a coverage-based OVM [open-verification-methodology] approach," says Vitesse's Chadra. The process started early in the design with behavioral models of the 24-port switch core and the MIPS CPU core to understand the dynamics of the chip under traffic flows. Verification then continued in increasing detail until the test bench was driving gate-level models with the clock-gating circuitry and the isolators in place. "We had specific targets in our verification plan based on our requirements document," Chadra says. "We augmented those targets with code-coverage metrics to guide the verification effort."

Redpine's Mattela says that the company's DVFS design required special care. Part of the problem was that logic simulators can't tell you whether a mismatch in signal levels would wreak havoc in a path between voltage islands. So Redpine verification engineers resorted to manual techniques, such as forcing nodes to tristate to see what would happen downstream. And part of the problem, Mattela warns, is that you never know the source of the models you are using. "Don't trust the models in a multivoltage situation," he states. "You don't know if they were written by an electrical engineer or by a software person who thinks a one is a one and a zero is a zero."

THE BACK-END FLOW

You now need to consider the physical-design phase: placement, routing, and design closure. During this phase, the impacts of IP reuse and design complexity begin to wane but do not by any means disappear. And the challenges of advanced processes cast a lengthening shadow over every step. First, the good news: Design managers seem to agree that tools have picked up and automated many of the new tasks that until recently had been manual. Madraswala says that Open-Silicon was able to take advantage of the DFM awareness of IC Compiler to help prepare for the complex design rules the process imposed. "A few years ago, everything about taking a power-managed design to tape-out was manual," says Mattela. "Now, there have been huge improvements, especially for postrouting validation."

The forces of change are still imposing problems, however. One is simply that new tasks breed new tools, and new tools are often problematic. "Let's say that some of the point tools are not mature," Chadra says. The tools' capacity is a more widespread issue. "We had to partition the design and run it through the tools in pieces," he explains. "Fortunately, most of the chip breaks into very natural segments. The biggest challenge was getting the switch through place and route."

Madraswala also cites place-and-route capacity. "You have a very limited design size when you switch on DFM awareness in IC Compiler," he says. "We were limited to about 400,000 placeable instances," yielding a small needle's eye through which to drive a 100 million-gate design.

Capacity was not the only issue with the place-and-route tools. Modern routers are timing-aware—that is, rather than trying to find the best possible route for every wire, the tools read the design's timing constraints and try to route to meet the timing on all nets. This process requires that the tool be able to estimate the delay on a proposed route, which in turn requires an estimate of the route's capacitance. So modern routing tools either call the sign-off extraction tool, which can be disablingly slow, or have built-in "quick-and-dirty" extraction estimators. Unfortunately, even at the 65-nm process node, parasitic extraction is a complex job for which there are no known quick approximations. "There are differences between IC Compiler and reality," Madraswala says.

Chadra is no more flattering. "The router capacitance estimates are not very accurate," he says, without specifying which place-and-route tool he is discussing. "We had the tool make some huge detours, which we had to go back and reroute."

The problem of timing estimation puts the EDA vendors in a dilemma. If the router's quick capacitance guesses are bad, physical-system designers will experience iterations between extraction, timing, and rerouting. Runtime and capacity will suffer if the router calls the sign-off extraction and timing tools, which have become complex because they must deal with all the effects of fine geometries.

After these chip designs were complete, both Cadence and Synopsys announced a third potential approach: moving preliminary placement and timing into the synthesis tools—even earlier in the design flow. It's not that the estimates will get any better but that the tool designers apparently hope to steer the synthesis tools away from creating netlists that the router will misestimate and misroute.

A similar issue exists with routers and design rules. If the router doesn't keep track of design rules as it works, many violations will emerge in the finished file. Routers thus pick up design rules from the LEF (layout-exchange-format) files and check the routes as they go. This process appears to have worked satisfactorily for digital circuits at the 65-nm node. Mentor Graphics' Madhani warns, though, that LEF can't express some of the rules, such as pinch rules, in advanced processes. Mentor thus now has its Olympus router dynamically call the Calibre sign-off tool for DRC on the fly. Again, this approach involves performance costs, but slow is better than wrong.

Perhaps surprisingly, after all the front-end work that has gone into them, power domains and third-party IP present issues for back-end design, as well. "Multiple power domains can lead to a complex closure," says Keh-Ching Huang, director of marketing at ASIC vendor Global Unichip. "We have to use a lot of manual procedures and scripts with them." Huang says that even IP choices influence the closure flow. "For instance, if a customer uses a low-speed DDR [double-data-rate] interface, the IP block usually comes in soft form and we have

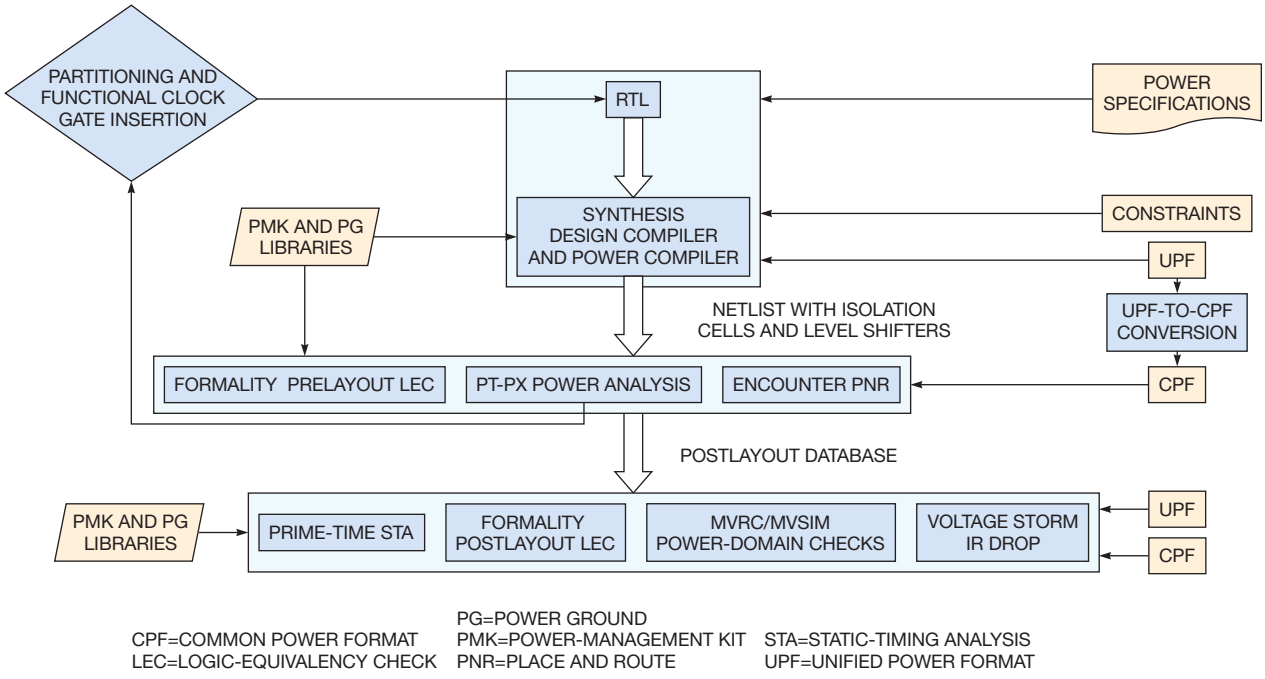

Figure 2 Redpine's methodology includes early capture of power intent and late checks on implementation.

to synthesize it. There will be timing-closure issues within the block. But if the customer licenses a high-speed DDR interface, it comes as hard IP, so the whole closure process is different. If there are issues, they will usually be with the package." Altogether, the impact of a design comprising mainly IP from outside sources on the struggle for design closure remains an underexplored question.

One final point is the impact of the new environment on analog design. Vitesse redesigned its copper PHY for this project, modifying the previous design to reduce power. In the process, the analog designers ran headlong into a number of layout-driven effects that were new in the 65-nm process. "We learned about well-proximity and drain-placement impacts on device performance," Chadra says. "The device models do an adequate job of modeling the effects, but we still had to do layout-extraction iterations to get the circuits performing as we desired."

So what is the big picture? Certainly SOC design today requires more upfront planning, especially for dealing with long routes, clocks, and power-management strategies. Verification planning upfront is also vital. Teams should understand that a lot is going on in synthesis tools. This step is no longer a straightforward replacement of Verilog statements by standard cells. Teams should thus plan to minimize iterations through the synthesis tools, especially once delicate structures, such as gated clock trees and test-scan chains, are in place. Likewise, teams should understand that aggressive power management can vastly complicate verification, and that this concern might justify choosing a more organic power-management strategy over a complex one.

Finally, physical design and closure are becoming more difficult. Choose front-end tools or develop scripts to head off congestion problems early. And plan for iterations between routing and the sign-off tools because they probably won't agree with each other. In basic structure, it's the same old flow. But the emphasis is shifting. "Probably 60% of the steps in this design were the same old steps," Madraswala says. "About 30 or 40% were specific to 65 nm, but those were the steps that caused the majority of the issues." EDN

FOR MORE INFORMATION

Brite Semiconductor
www.britesemi.com

Cadence
www.cadence.com

Global Unichip
www.globalunichip.com

HiSilicon
www.hisilicon.com

Mentor Graphics
www.mentor.com

MIPS
www.mips.com

Open-Silicon
www.open-silicon.com

Redpine Signals
www.redpinesignals.com

Synopsys
www.synopsys.com

Taiwan Semiconductor Manufacturing Co
www.tsmc.com

Virage Logic
www.viragelogic.com/index_en.asp

Vitesse Semiconductor
www.vitesse.com

You can reach Technical Editor Ron Wilson at 1-510-744-1263 and ronald.wilson@cancom.com.

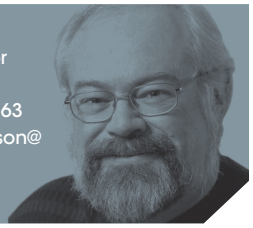

BY RAVI PODDAR, PhD, AND TED SUN • INTEGRATED DEVICE TECHNOLOGY INC

Power-grid analysis on SOC-graphics-chip design

YOU CAN SUCCESSFULLY RUN POWER-GRID ANALYSIS ON A 65-nm MIXED-SIGNAL SOC DESIGN.

Accurate power-grid analysis of 65-nm and smaller chips is becoming increasingly important to ensure reliable operation of devices in the field. Re-spins due to on-chip power-distribution issues are expensive and time-consuming and can lead to lost business opportunities. However, power-grid analysis of complex deep-submicron SOCs (systems on chips) with digital, analog, and third-party IP (intellectual-property) blocks can be a difficult task. The task becomes even more daunting when you couple it with looking at variations over multiple corners, as well as between static and dynamic analyses. This article explores a power-grid-analysis flow on a high-performance, 65-nm SOC-graphics-chip design and compares the results of different analysis types and corners.

The design is a 65-nm SOC with multiple IP blocks and memories from different vendors and a main clock frequency of 200 MHz. It has more than 4 million instances and multiple clock domains. Ensuring the reliability of the power-distribution network in this design is critical due to the mass-market nature of the product. This analysis addresses several challenges, including meeting an aggressive tape-out deadline, validating technology information, preparing data, handling analog and mixed-signal IP, resolving import and export issues among tools, examining IR (current/resistance) drop, performing EM (electromigration) analysis using static and dynamic methods, and identifying and repairing any issues the analysis uncovers. This analysis uses a typical flow (**Figure 1**).

INPUT-DATA PREPARATION

In the first step, you characterize power for standard cells and macros. This step generates power-characterized libraries by extracting the netlist and parasitic data from the cell's or macro's physical layout. You can do this characterization at different levels of detail, which, in turn, affect runtime and accuracy. You can obtain the power views for your IP in several ways. If your foundry officially supports your IP and power-analysis tool, there is a good chance that the IP creator has generated power views for your tool. If not, the vendor may be able to generate them for you. However, this process can take some time.

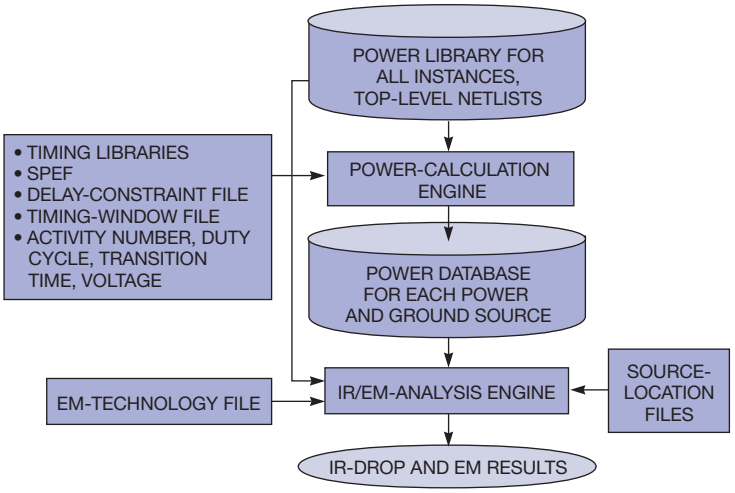

Figure 1 This design uses a typical power-analysis flow.

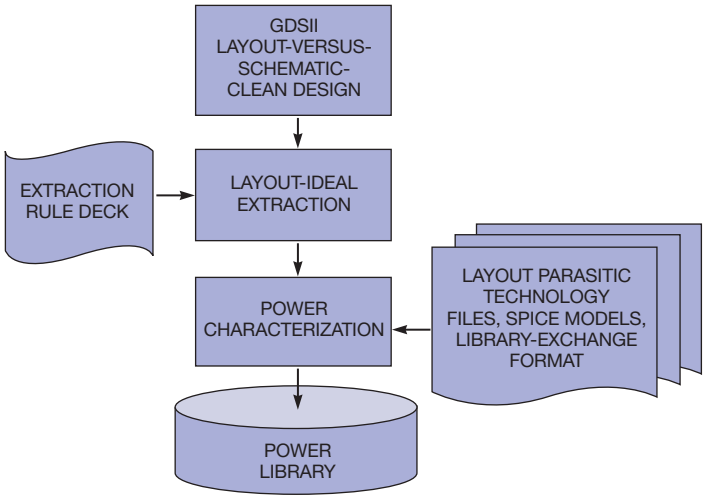

Figure 2 Generating custom power views can be nontrivial.

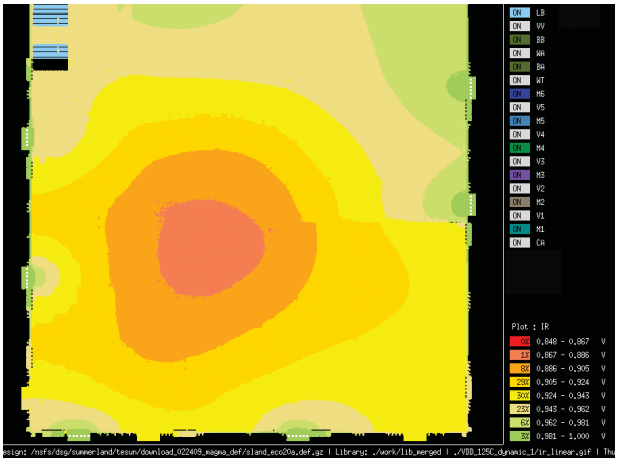

Figure 3 A display of the dynamic drain-to-drain-voltage analysis suggests an IR-drop issue.

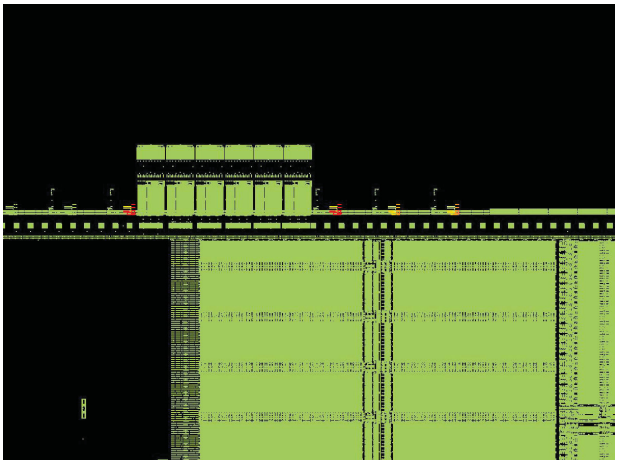

Figure 4 Zooming in on the data from Figure 3 pinpoints the problem.

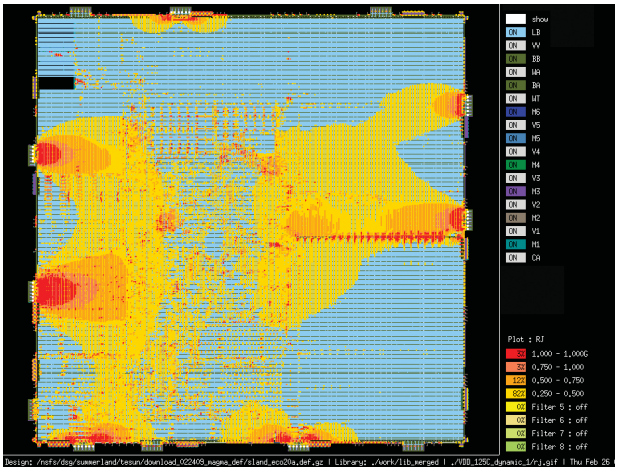

Figure 5 A current-density plot suggests the existence of some "hot" spots.

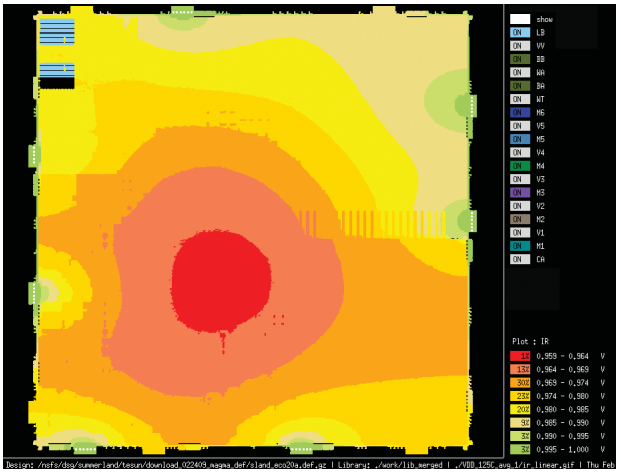

Figure 6 The static-IR-drop analysis does not suggest that problems occurred.

TABLE 1 POWER-CONSUMPTION VARIATION OVER DIFFERENT CORNERS (NORMALIZED TO 1W)

Corner	Internal power	External switching power	Leakage power	Total power
Fast	0.630	0.334	0.036	1
Typical	0.593	0.334	0.038	0.965
Slow	0.556	0.334	0.052	0.943

TABLE 2 IR AND EM RESULTS AT DIFFERENT ACTIVITY FACTORS

	V_{DD} IR drop (V)	V_{DD} EM	V_{SS} IR drop (mV)	V_{SS} EM
8%	0.984	20.448	14.496	25.267
10%	0.983	21.768	15.431	26.898
15%	0.981	25.08	17.767	30.984
20%	0.978	28.407	20.102	35.098

For the cases in which power views don't exist, including internally generated IP blocks, you must go through the custom power-analysis flow and generate the compatible views in that way (**Figure 2**). This task can be fairly complicated because you need to extract the layout, ensure that the LVS (layout versus schematic) is clean, and then run the power-characterization tool to analyze the block. Generating the correct technology files can be a difficult and time-consuming process, but you should again check with your vendors to see whether they have published the files. In this case, the foundry and the tool vendor generated all the technology files for the 65-nm process.

A number of other elements go into the power analysis, including the timing library, activity factors, power-source locations, physical netlist, extraction netlist, timing-window-constraint file, and SDC (standard-delay-constraint) file, all of which you must find or create. The timing library usually contains a power-look-up table that lists the gate's internal power. The power table is a function of input transition time and output net capacitance. The power-analysis tool uses

these entries to calculate overall power consumption within the cells. The activity factor—usually defined globally at the block level—represents how often nodes toggle in the design. It is also possible to have the activity factor propagate through the design logic and to set the factor explicitly on flip-flops to achieve a more realistic result. The most accurate method is to use a gate-level VCD (value-change-dump) file to determine exactly how the nodes toggle, but the onus is then on the user to find a VCD that represents the worst case for IR drop. Also, VCD-based analysis typically takes more runtime.

The power-source-location file specifies the exact X and Y locations for each power and ground source on the chip. This information enables the designer to justify whether the chosen positions will generate a balanced power distribution. At the chip level, these locations are usually on the power and ground pads.

The physical DEF (netlist-definition) file contains cell-placement, routing, and connectivity information for the SOC design. Although DEF is supposed to be a standard format, multiple issues can arise from some place-and-route tools during the generation of a complete DEF. Ultimately, you may have to write scripts to generate a complete DEF from multiple incomplete DEFs and use the new file to drive the power-analysis tools.

The extraction-netlist SPEF (standard-parasitic-exchange-format) file and the TWF (timing-window file) come from the normal process of design analysis. The wire parasitic data for resistance and capacitance are in the SPEF file, which the extraction tools generate. An STA (static-timing-analysis) tool generates the TWF. It contains information about switching windows and transition times, both of which have a large impact on dynamic power.

The SDC file contains clock-domain information, clock periods, and timing constraints. If the SDC file does not specify the default clock frequency, you can explicitly define it in the power-analysis tool's command file.

INSTANCE-BASED POWER CALCULATION

With all of the input files ready, you can begin the analysis, starting with instance-based power calculation. For each instance in the design, the power-analysis engine computes the power for each cell and reports it as an internal, an external, or a leakage type. Internal power is the power in the charging and discharging activities of a cell's internal capacitance, as well as the crowbar current. It is a function of load capacitance within the cell, voltage, frequency, and activity factors. In contrast, external switching power is essentially the power that is necessary for charging the wire capacitances that connect to the cell, which the input SPEF file determines. This power figure includes factors such as load capacitance, voltage, frequency, and activity factors. Leakage power is simply the static leakage power, independent of any cell activity.

Table 1 shows the design's power-calculation results at different corners. This analysis assumes propagated activity factors with maximum power normalized to 1W. As you would expect, you can see that a fast corner consumes the most power. Now, compare the results you obtain by using different values of parameters in the analysis command files.

TABLE 3 DYNAMIC AND STATIC ANALYSIS RESULTS OF USING GLOBAL AND PRIMARY INPUT ACTIVITIES

	Dynamic V_{DD} IR drop (V)	Dynamic V_{SS} IR drop (mV)	Static V_{DD} EM	Static V_{SS} EM	Total chip power (normalized)
Global activity (15%)	0.951	37.926	25.08	30.984	1W
Primary input activity	0.977	19.305	12.545	15.715	503 mW

Varying the global activity factor yields the results in **Table 2**. You can see when the global activity increased and that the IR drop and EM results are worse for drain-to-drain and source-to-source voltages. This expected change is not linear. After comparing the results at different global activities, you can also see what the difference would be between specifying global and propagated input activity factors. Setting the global activity factor to 15% and applying different activity factors for each pin yields the results in **Table 3**. Propagated activity factors show considerably reduced power consumption, IR drop, and EM violations.

Running static and vectorless dynamic analyses and using the most pessimistic results from each type to highlight issues in the design yields the results in **figures 3** through **6**. The worst IR drops occur while using dynamic analysis, and the worst EM violations occur in static analysis. Using fairly pessimistic operating conditions, including high-temperature and high-activity factors, accentuates errors. **Figures 3** through **6** highlight the dynamic drain-to-drain-voltage analysis and some weaknesses in the power grid at the top left of the die. These issues are not obvious in the static analysis. The resistor-current plot in **Figure 5** illustrates the current-density distribution in the design; the red areas are "hotter." After seeing these results, the physical-design team modifies the layout. After reanalysis, the design is ready for verification. Note that a similar analysis of source-to-source voltage is also necessary.

You can successfully run power-grid analysis on a 65-nm mixed analog and digital SOC design using industry tools and assistance from your foundry and the tool vendor. It is worthwhile to run static and vectorless dynamic analyses to find issues that one type of analysis may not highlight, particularly if you can obtain the input activity data from verification. EDN

AUTHORS' BIOGRAPHIES

Ravi Poddar is the director of CAD for the Design Automation Group at Integrated Device Technology. His main responsibilities include EDA-license management, custom-design methodology, physical verification, parasitic extraction, and simulation. Poddar has a doctorate in electrical engineering from the Georgia Institute of Technology (Atlanta).

Ted Sun is the staff CAD engineer in the Design Automation Group at Integrated Device Technology. His main responsibilities include physical verification, IR-drop and EM analysis, flow automation, and place-and-route support. Sun has a master's degree in engineering from the University of Dayton (Dayton, OH).

The Trade Shows that Connect You with the Electronics Manufacturing Industry

Electronics IPC Midwest

September 28–30, 2010
Donald E. Stephens Convention Center
Rosemont (Chicago), IL
ElectronicsMidwest.com

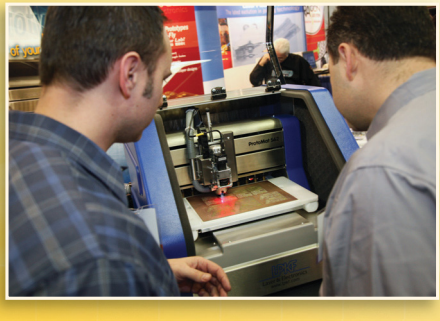

Electronics WEST

February 8–10, 2011
Anaheim Convention Center
Anaheim, CA
ElectronicsWestShow.com

Electronics NEW ENGLAND

April 6–7, 2011
Boston Convention & Exhibition Center
Boston, MA
Electronics-NE.com

Find the suppliers, tools, and services you need to make your product, process, and business more efficient, cost-effective, and profitable.

For information on attending or exhibiting, please call 310/445-4200 or visit CanonTradeShows.com

15512_AS_EL10

Produced and managed by Canon Communications LLC • 11444 W. Olympic Blvd. • Los Angeles, CA 90064-1549 • Tel: 310/445-4200

PERIODIC CHECKS OF YOUR TESTER'S RELIABILITY AND REPEATABILITY MINIMIZE THE CHANCE OF FALSE TEST RESULTS.

TEST YOUR TESTERS FOR R&R

BY MARTIN ROWE • SENIOR TECHNICAL EDITOR, *TEST & MEASUREMENT WORLD*

Automated test systems need to operate within designed parameters so that they produce reliable and repeatable results. False failures mean good products don't ship, and false passes can result in field problems. To keep a test system operating properly, you need to know how well it worked when it was put into operation or last calibrated, and you need to monitor the station for deviations.

The methods for testing a tester can vary depending on the tester's use and specifications, but there are many common techniques. One involves the use of check standards, which could be components such as resistors, voltage sources, current sources, or frequency sources that may be built into a test system or attached when needed. You can also use known-good devices—often called "golden" devices—to verify a tester's integrity, or you may want to rely on instrument self-test and loopback tests at regular intervals to keep testers running properly.

SET A BASELINE

Before putting a test system into service, you need to characterize it. Here's where a calibrated check standard or golden UUT (unit under test) can help. Such components let you establish a set of baseline measurements for each instrument and for the entire system. The best check standards are calibrated using methods traceable to NIST (National Institute of Standards and Technology) or another national laboratory. A golden UUT lets you check all of the connections in your system using the same cables, connectors, and fixtures that the system will use every day.

When you characterize a tester, always record and plot your results. You'll need them to compare against future measurements. Plots can reveal trends before the tester exceeds its limits. Over time, you'll make thousands of measurements that let you develop a performance history of R&R (reliability and repeatability).

Larry Raymond, president of Intrinsic Quality (www.intrinsic.com), a test-system developer, explains how to establish a baseline set of measurements for an in-circuit tester. "Use a reference board or, better yet, a set of boards with component values you trust," he says. Raymond recommends that you make enough measurements to establish a good statistical sample, typically hundreds to start. Then, you should plot a histogram so you can verify the tester's repeatability.

Raymond explains that one of his customers sometimes provides just one board for verifying a tester. When that scenario happens, he requires hundreds of measurements to prove that the tester is repeatable.

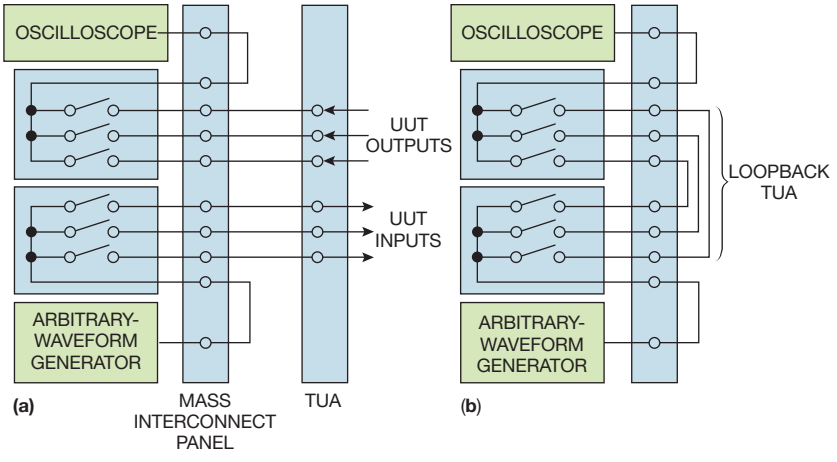

Figure 1 A TUA makes the instruments available to the DUT (a); a loopback TUA routes signals from the waveform generator to the oscilloscope for system test (b) (courtesy Ideal Aerosmith).

If possible, you should measure the components on several boards. That will give you a better statistical sample. It will also give you a clue about how the components' values vary, which lets you establish the tester's R&R.

If you have more than one sample UUT, you should look for inconsistencies in measurements. For example, you may find that a large set of capacitance measurements range from 120 to 132 µF. In such a case, Raymond suggests that you set the tester's acceptable limits for the component limits to perhaps 110 and 140 µF. If measurements fall outside those limits, you should check the tester against a reference board. Center your test window around the mean of the distribution of the baseline measurements. If the tester is operating within specifications, you may need to change tester parameters, such as test current, to compensate for the errors.

LOOPBACK TESTS

Letting the tester test itself is a common technique that many engineers use. Marius Gheorghe, engineering manager at test-system integrator Ideal Aerosmith (www.ideal-aerosmith.com), explains that the company often provides loopback TUAs (test-unit adapters) that connect signal sources to measuring instruments. **Figure 1** diagrams a test system that uses a loopback TUA to connect its oscilloscope and arbitrary-waveform generator.

DIAGNOSTIC PANELS CAN HOLD TEST COMPONENTS, SUCH AS RESISTORS, VOLTAGE SOURCES, CURRENT SOURCES, AND OSCILLATORS.

Figure 1a shows the normal operation in which the waveform generator's outputs and the oscilloscope's inputs come to a mass-termination panel. A TUA that attaches to the panel during normal operation routes signals to a DUT (device under test). Diagnostic software in the system controls both instruments and verifies that the waveform generator's signals are within limits.

Gheorghe suggests letting your loopback TUA make the connections between the signal sources and the measuring instruments and switching subsystems. "Don't route signals used for system self-test entirely through the system," he says. "Use an adapter. You'll get more flexibility."

Diagnostic panels such as the loopback TUA in **Figure 1b** may do more than just connect two instruments. They can hold test components, such as resistors, voltage sources, current sources, and oscillators. They can also connect instruments to any test and calibration resources that may reside in the test system, such as DMMs (digital

multimeters), oscilloscopes, and signal sources.

Test systems that consist mostly of card-based instruments may have just one stand-alone measuring instrument: an oscilloscope. Engineers often use the oscilloscope not only as part of the system, but also as a diagnostic tool. The oscilloscope can check signals as they pass through switches, cable, connectors, and mass-termination panels.

Mark Carlson spent 30 years as a test engineer with Texas Instruments (www.ti.com). During that time, he developed testers for the company's analog and digital ICs. In addition to using a loopback TUA, the testers he developed have a "calibration bus" that lets engineers check and calibrate power-supply outputs (**Figure 2**).

The calibration bus connects test equipment to resistors with 0.01% tolerance and to an ADC. The resistors provide a precision load for the power supplies, and the ADC measures power-supply voltages. Carlson explains that the ADC makes two voltage measurements on each power supply at the low and high ends of its range. A PC then calculates slope and offset to calibrate the power supply's output. The results go into a look-up table. Whenever the system needs an output voltage, it refers to the look-up table before sending a command to the power supply. The TUA also holds a NIST-traceable voltage reference for the ADC, making its measurements credible and verifiable.

Loopback tests and instrument calibration let you verify that a signal path is working properly, but what if the measurements seem wrong? You must troubleshoot the system. Here, reference devices can help you isolate errors.

For example, if your system has a DMM that measures power-supply voltage and the DMM has a self-test, then run the test. Try checking the meter with a known voltage that you've checked on another meter. Check the power supply's output. Is it what you expected? If it is, you likely have a problem with a cable, connector, or switch.

Test adapters that use golden UUTs may be far more complex than those that simply route signals between instruments. The board in **Figure 3** contains relays, an FPGA, and other devices that test engineer Todd Grey of Maxim Integrated Products uses to test digital ICs.

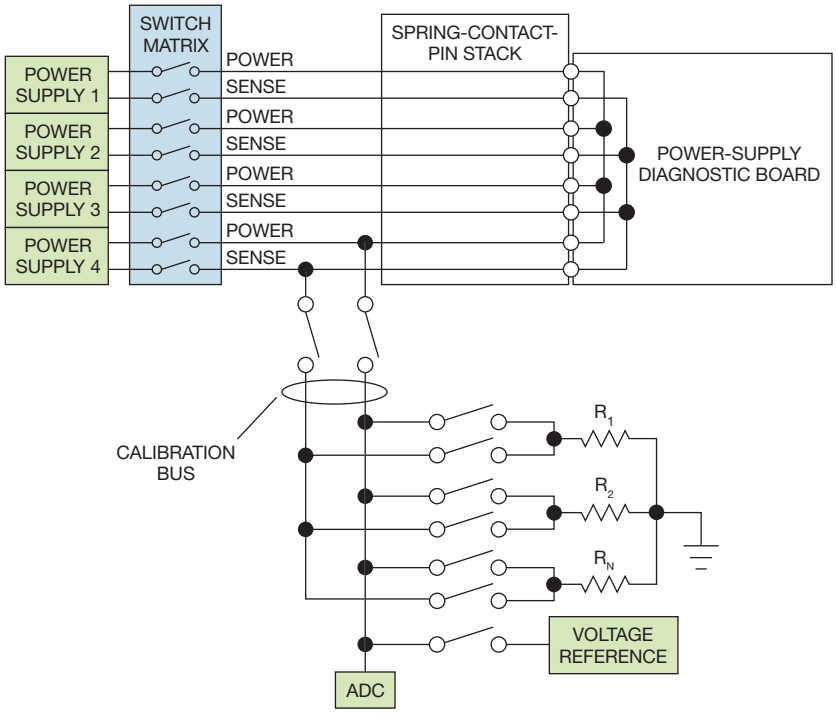

Figure 2 Switches connect a calibration bus to a test system's power supplies.

The FPGA works in combination with the relays under software control to connect a DUT to a power supply and to the appropriate pullup components for the DUT's communications bus. The pullups can be either resistors or FETs, depending on pullup requirements for the DUT.

HOW OFTEN TO CHECK?

Once you've established a baseline for a tester's performance, you should periodically check it. You can use a loopback setup, a reference component, or a golden UUT to conduct periodic checks. As a start, you can take advantage of self-test features built into your test equipment. You can run self-tests before every shift, every day, or every week, or you can run them at longer intervals, depending on how much you can trust the test equipment's stability.

Grey uses the test board in **Figure 3** at least once a week. He performs his tests using the tester's built-in diagnostics, a set of golden devices, and load boards. He will run complete characterization tests on the golden devices and compare the results against previous results. These tests provide greater test coverage than production tests. He often uses a separate DMM and oscilloscope when checking the tester.

Frequently checking your instruments can help you catch errors due to signal-path degradation or by an instru-

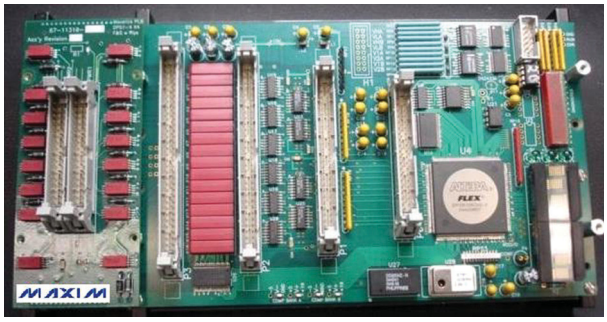

Figure 3 Load boards may contain connectors, relays, and electronics that engineers use to verify test-system integrity (courtesy Maxim Integrated Products).

LINEAR TECHNOLOGY

DESIGN NOTES

Dual Output High Efficiency Converter Produces 3.3V and 8.5V Outputs from a 9V to 60V Rail – Design Note 479

Victor Khasiev

Introduction

Among the many step-down DC/DC switching regulator controllers, the LTC3890 stands out because of its unique features. This 50µA quiescent current device can produce two output voltages ranging from 0.8V to 24V when powered from an input voltage of 4V to 60V.

Many high input voltage step-down DC/DC converter designs use a transformer-based topology or external high side drivers to operate from up to 60V_{IN}. Others use an intermediate bus converter requiring an additional power stage. However, the LTC3890 simplifies design, with its smaller solution size, reduced cost and shorter development time compared to other design alternatives.

Feature Rich

The LTC3890 is a high performance synchronous buck DC/DC controller with integrated N-channel MOSFET drivers. It uses a current mode architecture and operates from a phase-lockable fixed frequency from 50kHz to 900kHz. The device features up to 99% duty cycle capability for low voltage dropout applications, adjustable soft-start or voltage tracking and selectable continuous, pulse-skipping or Burst Mode® operation with a no-load quiescent current of only 50µA. These features, combined with a minimum on-time of just 95ns, make this controller an

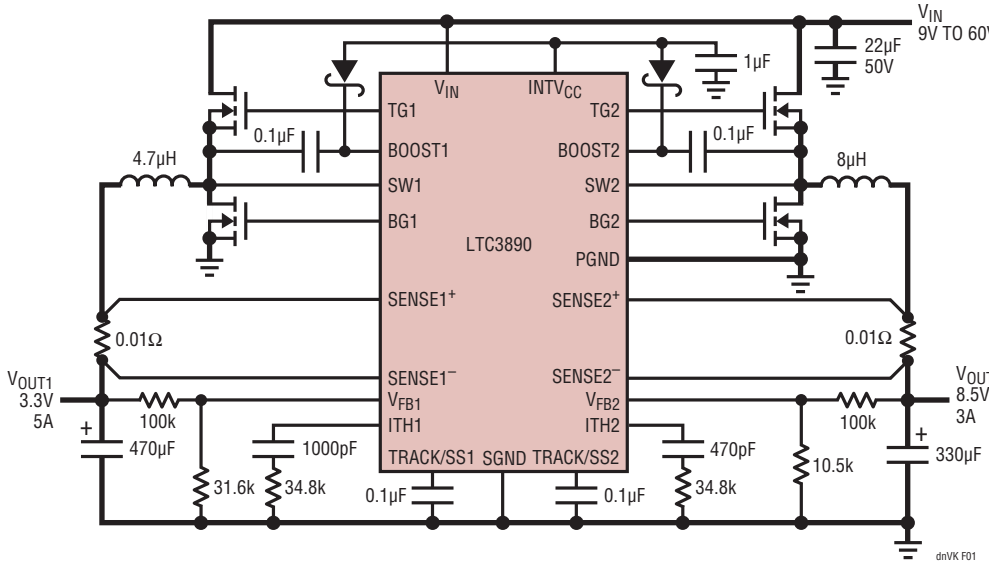

Figure 1. High Efficiency Dual 8.5V/3.3V Output Step-Down Converter

ideal choice for high step-down ratio applications. Power loss and supply noise are minimized by operating the two output stages out-of-phase.

Dual Output Application

Figure 1 shows the LTC3890 operating in an application that converts a 9V to 60V input into 3.5V/5A and 8.5V/3A outputs. The transient response for the 3.3V output with a 4A load step is less than 50mV (as shown in Figure 2).

Figure 3 shows the efficiency of the 8.5V channel with a 36V input voltage.

Single Output Application

The LTC3890 can also be configured as a 2-phase single output converter by simply connecting the two channels together. For example, a 9V to 60V input can be converted to an 8.5V output at 6A. Figure 4 shows the efficiency of this configuration at input voltages of 10V, 30V and 60V.

Current mode control provides good current balance between the phases. Less than 10% mismatch can be achieved, as shown in Figure 5.

Conclusion

Although there are many choices in dual-output controllers, the LTC3890 brings a new level of performance with its high voltage operation, high efficiency conversion and ease of design.

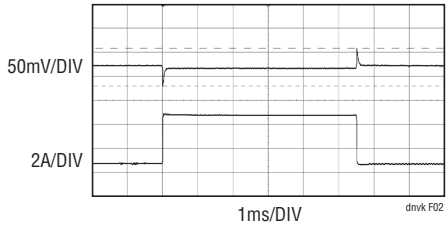

Figure 2. Transient Response of 3.3V Channel (I_{OUT1}: 1A to 5A)

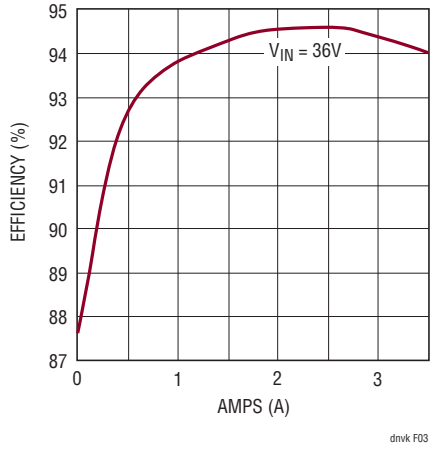

Figure 3. Efficiency of the Converter in Figure 1 for the V_{OUT2A} 8.5V Channel

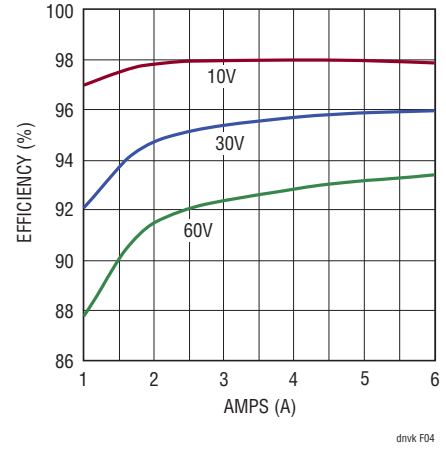

Figure 4. Efficiency of the LTC3890 Configured as a 2-Phase Single Output of 8.5V at Up to 6A from a 10V to 60V Input

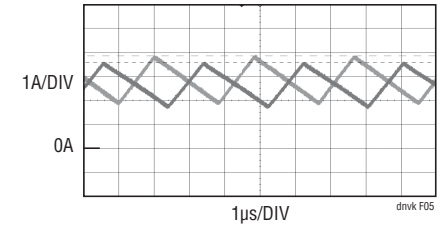

Figure 5. The Inductor Current in a 2-Phase Single Output Converter. Currents in Both Inductors Shown with a 24V Input and 8.5V at 6A Output.

Data Sheet Download

www.linear.com

For applications help,
call (408) 432-1900, Ext. 3161

Linear Technology Corporation
1630 McCarthy Blvd., Milpitas, CA 95035-7417
(408) 432-1900 • FAX: (408) 434-0507 • www.linear.com

dn479f LT/TP 0610 116K • PRINTED IN THE USA

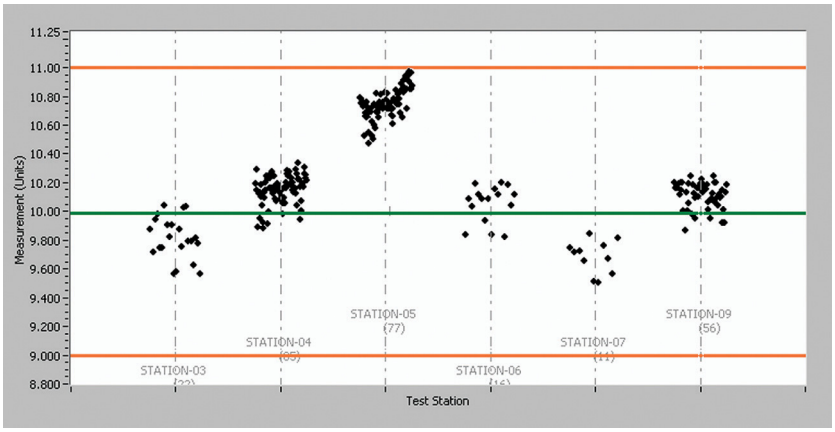

Figure 4 Test data from six test stations shows that readings from Station 5 are near the high limit (courtesy Tektronix).

ment's drifting out of calibration, especially if you use that instrument to test an important product specification. "The last thing you want is to have a calibration lab report that one or more specifications for an item of test equipment was received out of tolerance," says David Buxton, senior test engineer at Tektronix (www.tektronix.com). "When that [situation] happens, an investigation should follow, which can lead to a recall notice for a free recalibration of the instrument."

Buxton notes that you can apply SPC (statistical-process-control) analysis to the test data to look for trends. If the system has a built-in reference component, then you can program it to automatically run a check and log the results.

Checking production test data can provide clues about a tester's health. Look for trends in the measurements on production parts. A trend could indicate that an instrument is drifting out of its tolerance or that something else is wrong. **Figure 4** gives an example of measurements made on five test stations, and **Figure 5** shows data from a single tester.

In **Figure 4**, the data from Station 5 indicates a measurement problem with the station. For example, if you notice that a power supply's voltage is dropping, it could indicate an increased resistance in a switch, connector, pin, or relay. Test pins get soft, and they may not make good contact over time. In addition, poor test-clamp alignment may also cause measurement errors.

Figure 5's plot highlights a drop in a measured value, and the measurements are erratic. This drop could indicate an intermittent instrument problem or a signal-channel problem that will require troubleshooting.

If you encounter situations such as those in **figures 4** and **5**, you can start by checking the test equipment against a known reference. You can also run the same measurement on a different instrument or tester. If you see significant differences in a measurement from one tester to another, you should suspect the tester.

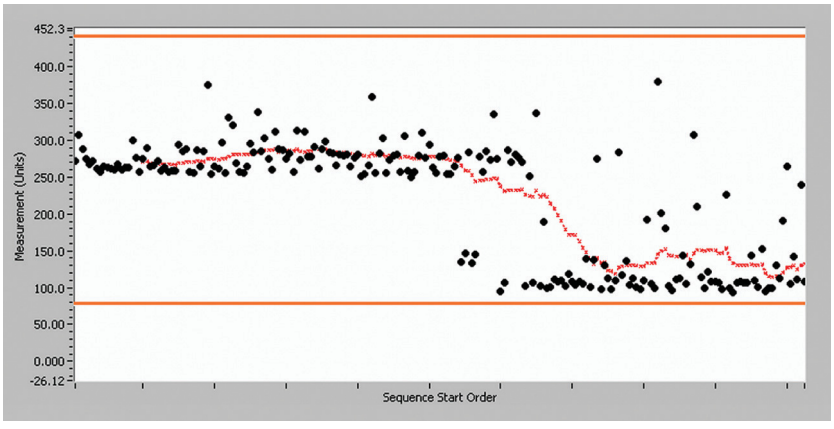

Figure 5 Test data shows a sudden measurement dropoff, but not from all measurements, which could indicate intermittent performance (courtesy Tektronix).

You should try to minimize the number of obstacles in the suspected measurement path. Running a signal from a signal generator through an amplifier, an attenuator, and a test head to a spectrum analyzer means that any of them could cause a measurement error. You will likely have to check each component.

Signal paths that carry RF signals such as radar or serial data streams of 10-GbE (gigabit/sec Ethernet) have more possibilities for signal errors than low-frequency signals because of losses and reflections. Chris Scholz, field applications engineer at LeCroy, recommends that you use a TDR (time-domain reflectometer) to characterize signal paths and that you periodically rerun TDR measurements, looking for changes in results that could indicate a change in impedance.

TDRs are usually options on high-bandwidth sampling oscilloscopes. Their wideband analog front ends let them measure reflections on repetitive pulses with rise times of tens of picoseconds (**Reference 1**). You can use TDR measurements to calculate S parameters that characterize a signal path in the frequency domain. You can also use a VNA (vector network analyzer).

Knowing the characteristics of your signal path, you can "de-embed" or compensate for channel losses in your measurements. If you shut down your tester for maintenance every six months, for example, then that's the time to run a TDR measurement. You need to make measurements with a TDR or VNA for every test fixture that carries high-frequency signals. EDN

REFERENCE

1 Rako, Paul, "TDR: taking the pulse of signal integrity," *EDN*, Sept 3, 2007. www.edn.com/article/462661-TDR_taking_the_pulse_of_signal_integrity.php.

A version of this article appeared in the April 2010 issue of *EDN*'s sister publication *Test & Measurement World*.

USER-CONFIGURABLE FPGA MODULES BOOST PXI-SYSTEM VERSATILITY

FPGA PERFORMANCE AND CONFIGURABILITY ARE BECOMING AVAILABLE TO SYSTEM DEVELOPERS AND END USERS, ADDING A NEW RANGE OF DESIGN POSSIBILITIES FOR PXI SYSTEMS.

BY RICHARD A QUINNELL • CONTRIBUTING TECHNICAL EDITOR

The advent of PXIe (Peripheral Component Interconnect Express) created new opportunities for PXI systems to handle more data faster. Module developers turned to FPGAs (field-programmable gate arrays) as hardware-configurable alternatives to high-speed or multicore CPUs for handling this data. Now, such FPGA performance and configurability are becoming available to system developers and end users, adding a new range of design possibilities for PXI systems.

PXI-module developers began using FPGAs in their high-performance designs nearly a decade ago to handle ever-increasing data rates and to reduce latency. The configurable hardware offered significant performance advantages over software-based designs. Sebastien Maury, Sundance Multiprocessor Technology's (www.sundance.com) regional director for the Americas, estimates that an FPGA can provide 20 to 30 times more digital-processing performance than a PXI host controller.

According to Ryan Verret, FPGA-for-test product manager at National Instruments (www.ni.com), the performance boost can be as much as several orders of magnitude. "It's really staggering what computation you can do on an FPGA," he says. "Some devices have more than 500 DSP blocks on them, allowing you to do lots of FFTs [fast Fourier transforms] in real time."

Flexibility is another reason for incorporating FPGAs into module designs. David Manor, vice president of hardware engineering at Geotest-Marvin Test Systems (www.geotestinc.com), points out several advantages that stem from field configurability. "Using FPGAs allows us to get products that relate to new standards out early, even before the standards are fully defined," he says. "We have to make some assumptions, so we may not get everything exactly right at first, but [we] can release updates and bug fixes for users to download without the need to return their boards to us."

The use of FPGAs also simplifies the creation of unique functions for customers. "The customer simply installs new firmware to gain new features," Manor says. He adds that the ability to upgrade and modify module functions in the field helps keep Geotest's products viable longer in a given application.

CONFIGURABLE FPGAs

Until recently, however, the FPGAs in PXI modules have been nearly inaccessible to end users. Upgrades and enhancements came from module vendors, and users could not readily implement their own design ideas. That situation changed in the last year with the introduction of user-configurable FPGA modules for PXI from at least four manufacturers: Geotest, NI, OpenATE (www.openate.com), and Sundance.

The Geotest GX3500 PXI module connects 160 digital I/O channels to an Altera (www.altera.com) Cyclone III FPGA with 55,000

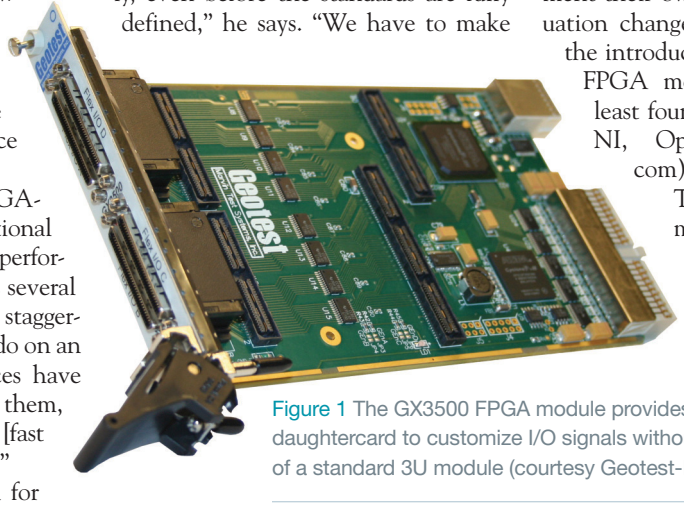

Figure 1 The GX3500 FPGA module provides space for an internal daughtercard to customize I/O signals without exceeding the envelope of a standard 3U module (courtesy Geotest-Marvin Test Systems).

logic elements, four PLLs (phase-locked loops), and 2.34 Mbits of memory. The FPGA has access to all PXI-bus resources, including clocks and triggering. The module also supports an internal expansion-card assembly that customers can use to customize the front-panel I/O connections (**Figure 1**).

NI offers the FlexRIO family of FPGA modules, which uses Xilinx (www.xilinx.com) Virtex-5 FPGAs of varying capacities. The NI PXI-795x series comprises conventional PXI modules, whereas the PXIe-796x-series instruments are PXIe modules. FlexRIO modules offer 132 I/O lines and accept front-panel-mounted adapters that customize the I/O interface for connection to analog, Ethernet, IEEE 1394 (FireWire), Camera Link, and other specialized interfaces (**Figure 2**). Adapter modules are available from NI as well as from third-party partners, such as Adsys Controls, Averna, NexFrontier Solutions, and Prevas (www.adsyscontrols.com, www.averna.com, www.nexfs.com, www.prevas.com).

OpenATE also uses a Xilinx Virtex-5 FPGA on its FPGA carrier card. The OpenATE card offers a more basic design than those from Geotest and NI, however, with the FPGA handling the PXI-bus interface. The company does provide PXI-interface IP (intellectual property), as well as IP for a DIMM (dual-inline-memory-module) interface. The FPGA carrier supports a user-defined daughtercard that handles 156 I/O lines but provides no built-in front-panel connections. Users must customize the front panel along with the daughtercard.

The Sundance SMT-700 FPGA card offers a PXIe interface, a variety of front-panel serial interfaces, and internal digital-I/O headers along with a choice of several Xilinx Virtex-5 devices. Front-panel serial interfaces include 10/100/1000 GbE (gigabit Ethernet), fiber optic, and USB (Universal Serial Bus). Dual internal mezzanine connections allow developers to attach Sundance analog modules or additional digital I/O to the FPGA card. The connections reside on both sides of the board, causing the FPGA card to occupy two PXI slots when populated.

CUSTOMIZED I/O

Although these FPGA modules have significant differences, they do share some attributes. For example, you can convert each of them from a simple data-processing card into a fully defined instrument with the addition of customization circuitry, through either a mezzanine or an extension card. FPGA-module vendors offer both predefined cards and open specifications from which users can develop their own cards. This ability to define the module's I/O-signal conditioning and formatting combines with the FPGA's configurability to give users an unprecedented opportunity for creating innovative PXI-module functions from off-the-shelf building blocks.

An FPGA module with high-speed ADCs on the customization card, for instance, has all the necessary hardware to serve as a digital spectrum analyzer. Because of the tremendous parallelism available in the FPGA, the analyzer can provide continuous monitoring across a wide frequency band in place of the usual swept-spectrum analysis. As NI's Verret points out, this feature allows the FPGA-based instrument to analyze time-multiplexed communications protocols, such as RFID (radio-frequency identification), that use short energy bursts. Swept-spectrum instruments can easily miss these bursts. Verret also notes that the continuous monitoring allows an FPGA-based instrument to generate triggering signals based on complex power-frequency spectrum masks. This feature can be helpful in reducing the capture-depth requirements of downstream data-acquisition systems.

Developers can also configure an FPGA module to dynamically generate test signals, reducing the number of test vectors necessary in an ATE (automated-test-equipment) system. For example, when a conventional ATE system needs to send data to a communications port on an IC under test, the system must use a long series of test vectors to drive arbitrary-waveform generators that produce the signals driving the port. These vectors must describe the signal values at every time

step throughout the data transfer and can be long. By using an FPGA, developers can create a state machine that can control the signal timing so that the test vectors describe only those data values that will be transferred and not the signal's entire time history.

Dynamic signal generation also allows an FPGA-based instrument to handle situations that are difficult or impossible to handle with static vectors. Testing of RFID devices, for instance, is difficult with static vectors because they involve multiple stimulus-response interactions, which may have variable timing. Similarly, the testing of an engine-control module requires that the tester emulate how the engine and drive shaft will behave in response to the control signals. Such responses are virtually impossible to emulate with vectors, and host-controller-based emulation is often too slow or has too much latency. An FPGA-based emulator can resolve all these issues.

FPGA modules can even adapt their functions on the fly. According to Sundance's Maury, systems can partially reconfigure an FPGA while the module is running to change its behavior in response to incoming signals. Such dynamic changes might include altering test vectors to more extensively test functions that appear marginal or automatically revising signal-processing-algorithm parameters to tune them for the current test conditions.

DEVELOPMENT SUPPORT

The extreme flexibility that user-configurable FPGA modules provide PXI developers does come at a price. NI's Verret, for instance, points out that designing the logic of an FPGA may not be within the abilities of many test engineers because of the need to work with an HDL (hardware-description language) to define their designs. "Even if they have the skills," says Verret, "[configuring an FPGA] is harder to do than writing processor code." To minimize the FPGA-configuration effort, most module vendors offer development support. NI, for instance, has created an FPGA-module extension for its LabView development tool. This extension allows developers to describe their entire test-system design and data-processing algorithms in LabView and have the tool automatically prepare the FPGA-programming code. The company is also providing a library of IP for more complex, application-specific functions. NI's third-party partners offer additional IP as well as full turnkey solutions.

Figure 2 The FlexRIO-series FPGA modules use external adapter modules for I/O customization and are available in PXI and PXIe versions (courtesy National Instruments).

Geotest has made its FPGA module fully compatible with standard Altera object files, meaning that developers can use the free Quartus II Web Edition development tools that Altera provides in support of its FPGAs rather than needing proprietary tool sets. Geotest also provides drivers and a virtual panel for interactive design, debugging, and deployment of the FPGA configuration. Interface files for this tool support programming tools and languages, such as ATEasy, C/C++, Visual Basic, and LabView. The company offers an online tutorial to help developers get started.

Sundance, which caters to more experienced hardware-design customers, provides only some basic software to developers. For full development support, the company works in partnership with software company 3L (www.3l.com). The 3L Diamond tool set allows developers to create a task-based model of the FPGA's function, and it then automates the design's compilation to FPGA-programming vectors. It also supports the synchronization of multiple FPGAs for achieving extended performance through multiprocessing. Developers can also use a tool such as The MathWorks (www.mathworks.com) Matlab to capture their designs and convert the tool's output to an HDL for compilation using Xilinx tools.

WILL AN FPGA HELP?

Although the level of programming involved in customizing a user-configurable FPGA module for their applications may be daunting to many PXI users, the results can be well worth the effort. The key is deciding whether the effort is necessary. NI's Verret recommends that users consider an FPGA module if they can benefit from custom triggering that will reduce the amount of data the host controller must process and when data-processing requirements demand the performance that an FPGA provides. Geotest's Manor says that the use of an FPGA card is justified when there is nothing available to support a special function that an application requires or when there is a need to change functions on the fly.

Ultimately, customer feedback will determine the long-term future of these user-configurable PXI modules. Vendors anticipate continuing to offer larger, faster FPGAs as they become available and will likely create additional daughtercards and IP based on the frequency of customer requests.

Even if the FPGA module remains only a specialty item, however, the future is secure for FPGAs in stock PXI-instrument modules. "We couldn't be where we are today in PXI without FPGAs," says Geotest's Manor. "They have been a key component in the success of PXI as a performance platform." EDN

A version of this article appeared in the May 2010 issue of *EDN*'s sister publication *Test & Measurement World*.

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Isolated FET pulse driver increases power rate and duty cycle

Jaime Castelló, José M Espí, and Rafael García-Gil
University of Valencia, Spain

In power converters, pulse-drive circuits transmit the pulses a controller generates to the power transistor. Driver circuits must both transmit the controller's switching on/off signals with galvanic isolation and provide energy to turn the switch on and off and to maintain the required on or off state. The required energy increases with the power transistor's input capacitance, which also increases with the power that the transistor module manages. Thus, when the circuit requires high power, designers typically parallel the power transistors, increasing the input capacitance. When you need to operate IGBT (insulated-gate-bipolar-transistor) modules in parallel, it is best to share the gate drive because using different driver circuits introduces additional variation in turn-on and -off times and creates a possible imbalance between each power module.

The basis of the circuit in **Figure 1** is an earlier Design Idea (**Reference 1**). The operation of the circuits is basically the same as the one in the previous Design Idea, but this one can drive MOS-

DIs Inside

▶To see all of *EDN*'s Design Ideas, visit www.edn.com/designideas.

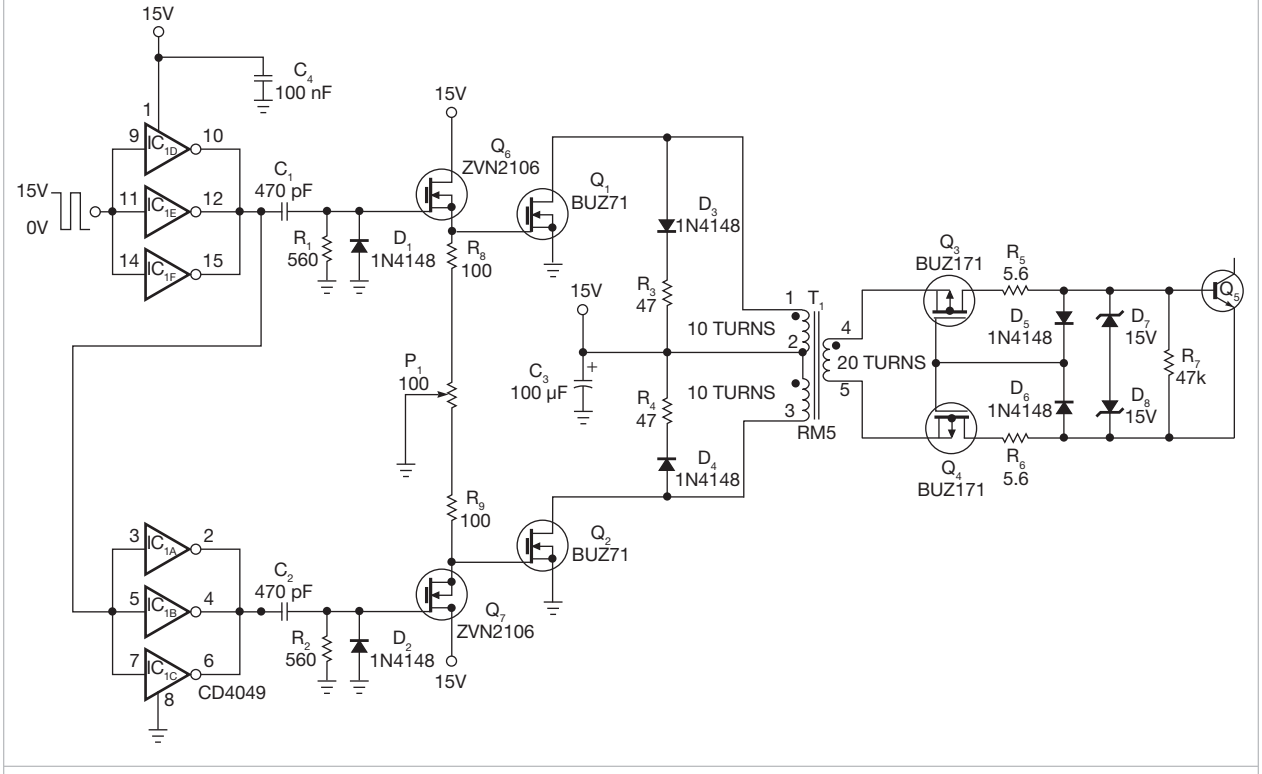

Figure 1 This isolated pulse driver can transmit all duty cycles, even with high-power MOS-FET/IGBT modules that have large input-gate capacitance.

Energy Harvesting Now

Free Power from Thermal, Kinetic & Solar Energy

Our new analog IC solutions enable the commercial deployment of energy harvesting from a variety of "free" energy sources. An appropriate transducer placed on the energy source delivers an electrical signal that our products convert and condition into usable power. These revolutionary ICs consume only nanoamps of current to provide high efficiency power conversion with minimal external components.

Energy Harvesting IC Family

Part Number	Description	Energy Source
LTC®3105	400mA boost converter with MPP control and 250mV start-up	
LTC3108	Ultralow voltage boost converter and system manager	
LTC3109	Auto-polarity version of LTC3108	
LTC3588	Piezoelectric energy harvesting power supply	
LT®3652/HV	Power tracking 2A solar battery charger	
LTC4070	Nanoamp operating current shunt Li-Ion battery charger	

Info & Free Samples

www.linear.com/energyharvesting

1-800-4-LINEAR

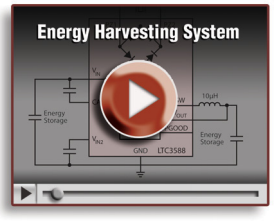

LINEAR TECHNOLOGY

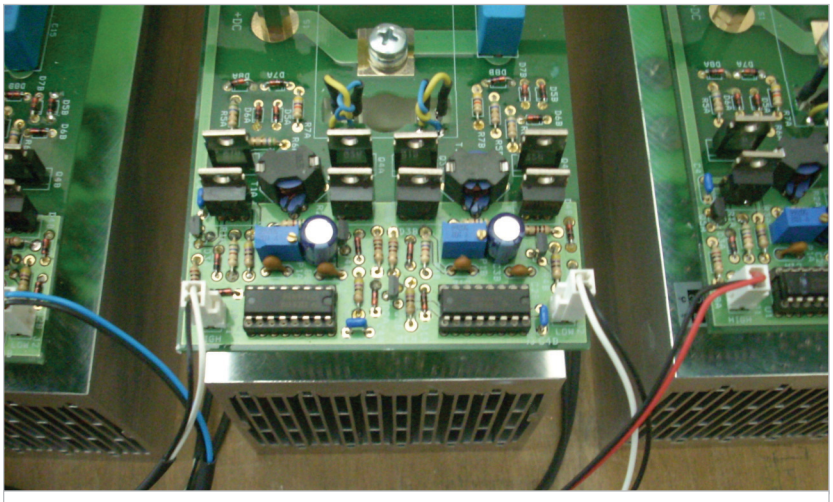

Figure 2 An isolated FET pulse driver in a 10-kW, three-phase inverter for grid injection requires few components and has galvanic isolation.

FETs or IGBTs with input capacitances higher than 5 nF. This circuit provides full galvanic isolation and requires no floating power supplies; it can transmit duty cycles that approach 100%.

This circuit adds transistors Q_6 and Q_7 to the circuit in the previous Design Idea. Transistors Q_1, Q_2, Q_3, and Q_4 are now higher power because they can manage a larger current depending on the transistor they need to drive. Transistors Q_1 and Q_2 are BUZ71 units, Q_3 and Q_4 are BUZ171 devices, and Q_6 and Q_7 are ZNV2106s. The differentiator circuits, C_1/R_1 and C_2/R_2, generate 1-µsec-long pulses, and you do not apply them directly to the gates of the Q_1 and Q_2 transistors, as in **Reference 1**, but to transistors Q_6 and Q_7. Although the input capacitances of Q_1 and Q_2 are nearly 700 pF, the input capacitances of Q_6 and Q_7 are approximately 75 pF, ensuring that the narrow pulses will transmit properly.

During the rising edge of the drive-control signal, Q_7 turns on, and its current starts charging the input capacitance of Q_2 through the on-resistance of Q_7. Because Q_7's on-resistance is only a few ohms and no additional drain resistance exists, the charging process of Q_2's input capacitance becomes fast, although its input capacitance is high.

As the gate voltage of Q_2 increases, the gate-to-source voltage of Q_7 decreases, and the transistor turns off. As a result, the narrow pulses that the differentiator circuit generates transmit to transistors Q_7 and Q_2 through coupling transformer T_1 to transistor Q_3, which charges Q_5's gate-to-source input capacitance. The same process occurs with Q_6, Q_1, and Q_4 during the falling edge of the drive-control signal to discharge Q_5's gate-to-source input capacitance.

With potentiometer P_1, you can control the discharge time of Q_1 and Q_2 and thus adjust the offset of the drive signal you apply to the power transistor. Because Q_6/Q_1 and Q_7/Q_2 transmit narrow pulses and have fast rising and falling edges, you can obtain a great duty-cycle variation even for high switching frequencies. You can control the duty cycle from 2 to 98% with a 20-kHz switching frequency. The circuit's compact design lets you mount it close to the power module, which minimizes parasitic elements.

Figure 2 shows the driver prototype for a 10-kW/20-kHz three-phase power inverter for grid injection. The circuit uses SKM75GB128 power transistors from Semikron (www.semikron.com). The transistors have a measured input capacitance higher than 15 nF. In this situation, the total current consumption of the FET pulse driver is lower than 30 mA. EDN

REFERENCE

1 Espí, José M, Rafael García-Gil, and Jaime Castelló, "Isolated FET pulse driver reduces size and power consumption," *EDN*, March 30, 2006, pg 98.

Detect missing pulses to avoid losing data

Chien-Hung Chen and Tai-Shan Liao,
National Applied Research Laboratories, Hsinchu, Taiwan

THIS CIRCUIT SOLVES THE PROBLEM OF A DATA-CONVERSION SYSTEM THAT DOES NOT ENABLE AN ACQUISITION IN THE FIRST OF 256 PULSES.

In an SWIR (short-wave-infrared) indium-gallium-arsenide-based imaging system, the system must detect every trigger pulse; otherwise, lost data will result. This Design Idea solves the problem of a data-conversion system that does not enable an acquisition in the first of 256 trigger pulses. **Figure 1** shows the SWIR sensor, IC_1. The AD-Trig signal, Pin 24, produces one pulse for each of 256 pixels. **Figure 2** shows the AD-Trig pulses, Channel 1 (Trace 1) and Channel 2 (Trace 2), as the two pulse trains. Channel 3 (Trace 3) is the electro-exposure time signal. The signal is active-high when you expose the image sensor to light. The signal goes low before the AD-Trig signals begin.

Oscilloscope Channel 4 (Trace 4) from the microcontroller's I/O pin goes high during the second pulse, but it should go high during the first pulse. Thus, the microcontroller missed the first pulse and the first pixel of the sensor image. A hardware and firmware redesign solves the problem. The AD-Trig signals connect to a microcontroller's PA0 pin, which you can use as an external interrupt input. **Figure 3** shows the program flow, and **Listing 1**, which is available at www.edn.com/100610dia, shows the source code.

In **Figure 3**, the left side is the TD-poll routine in the **listing**. It uses an I/O pin to detect the electro-exposure signal. When the electro-exposure time signal finishes, the firmware enters the external interrupt subprogram when you activate the AD-Trig signal. The interrupt subprogram reads all of 256 pixels data. The pixel data is no longer lost because the time for firmware processing of the TD poll differs from that of the external interrupt subroutines, and the external interrupt has a higher priority than does the TD poll. EDN

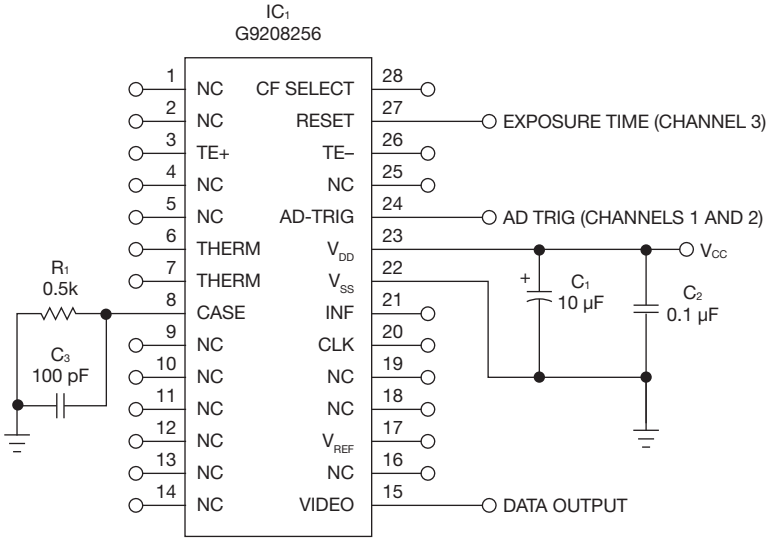

Figure 1 An SWIR sensor's AD-Trig signals produce one pulse per pixel.

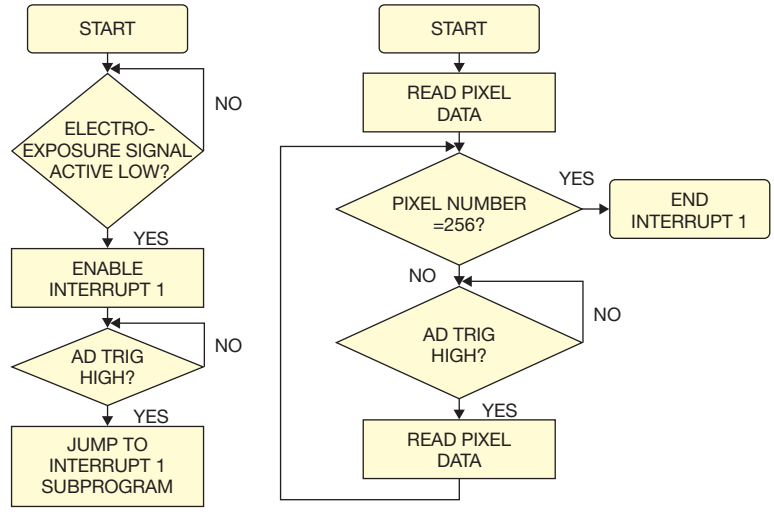

Figure 3 Using an interrupt, the firmware code detects the first pulse, AD-Trig high.

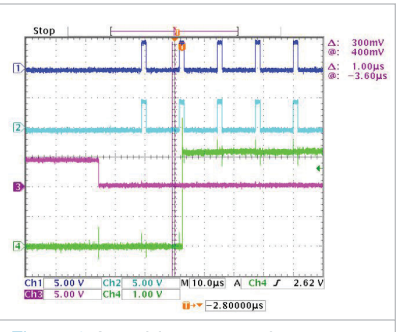

Figure 2 A problem occurs because the microcontroller missed the first pulse on Channel 4.

Circuit lets you measure zener voltages and test LEDs

Vladimir Oleynik, Moscow, Russia

To measure a zener diode's breakdown voltage, you need a dc voltage source whose voltage exceeds that of the zener voltage. In **Figure 1**, resistor R_{SER} provides voltage drop between V_{IN} and V_{ZEN}. In any case, V_{IN} should exceed V_{ZEN}. Resistor R_{SER} must provide current, I_{ZEN}, that can keep the zener diode in reverse breakdown. That is, the current must be more than I_{ZEN}–I_{ZENMIN} and less than I_{ZEN}–I_{ZENMAX}. You also need to consider the current that flows through the load. Otherwise, V_{ZEN} will be unregulated and less than the nominal breakdown voltage. Also, the power that the zener diode dissipates should not exceed the manufacturer's specifications. Except for the value of I_{ZEN}–I_{ZENMIN}, all necessary data appears in zener-diode data sheets.

The circuit in **Figure 2** uses one or two AA/AAA cells, which ensures testing irrespective of the value of the tested zener voltage to approximately 20 to 25V. The heart of the circuit is a Zetex (www.diodes.com) LED driver, ZXLD381. It operates mainly from 1.5 or 1.2V cells, and it has a maximum input voltage of 10V. The LED driver generates constant-power pulses that charge output of the 10-µF capacitor, which requires low leakage current. The capacitor's voltage provides constant current through R_1 and tthe zener diode that connects in series. When you connect output probes to a digital multimeter's V and COM sockets, you can directly measure the zener's voltage when the S_2 switch is in the position **Figure 2** shows.

When S_2 is in the upper position, the meter measures voltage across R_1, a 1-kΩ resistor; the meter displays a negative-voltage-drop value. R_1's value ensures direct reading of the meter; the voltage drop across R_1 corresponds to the zener's current, so there is no need to switch over DVM (digital-voltmeter) ranges. R_1's voltage drop limits the zener diode's voltage value that the circuit measures. If R_1's value is 1Ω, then the voltage drop across it is insignificant and, in millivolts, is equal to the zener diode's current in milliamps.

If you need to measure zener-diode voltage higher than 20 to 25V, you can add an LED driver (**Figure 3**). When the S_2 switch is in the upper position, both LED drivers connect in series, and you can measure zener voltage to approximately 40V at 0.7 mA. With S_2 in the lower position, both LED drivers connect in parallel, and the tested zener voltage is approximately 20 to 25V at several milliamps of current. Both parallel and in-series connections provide zener-diode-voltage measurement at two current values. In some cases, I_{ZEN} values may not fit the **equation** I_{ZEN}–I_{ZENMIN}<I_{ZEN}<I_{ZEN}–I_{ZENMAX}, and the zener diode may go out of regulation.

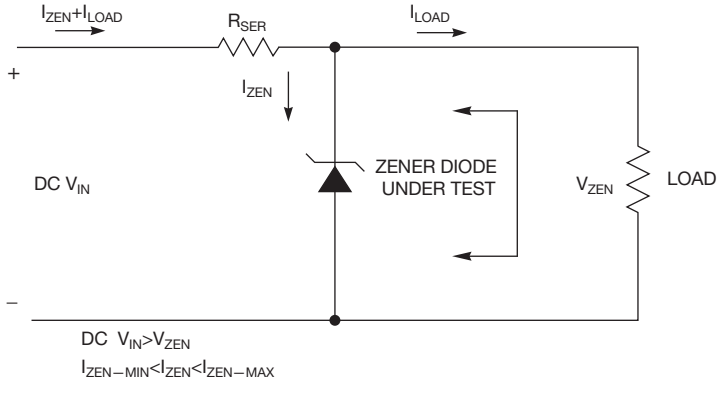

Figure 1 Current through a zener diode creates a fixed voltage across the device.

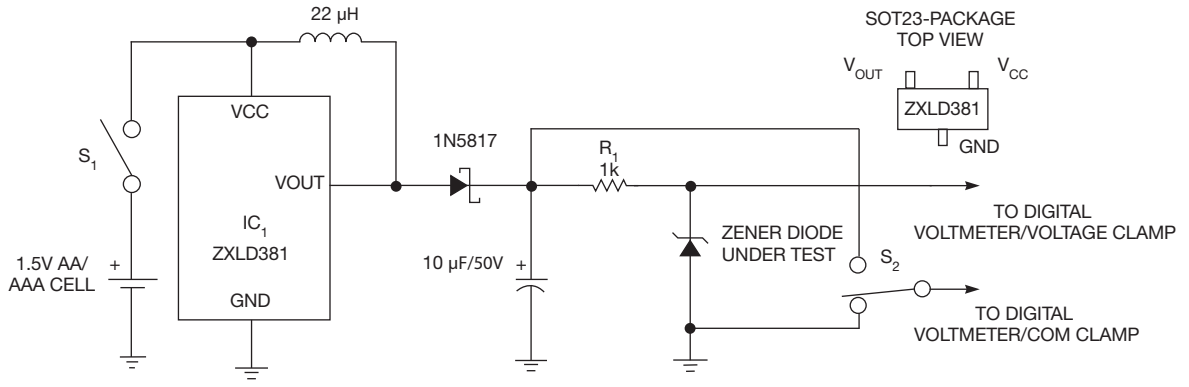

Figure 2 An LED driver provides current for the LED under test.

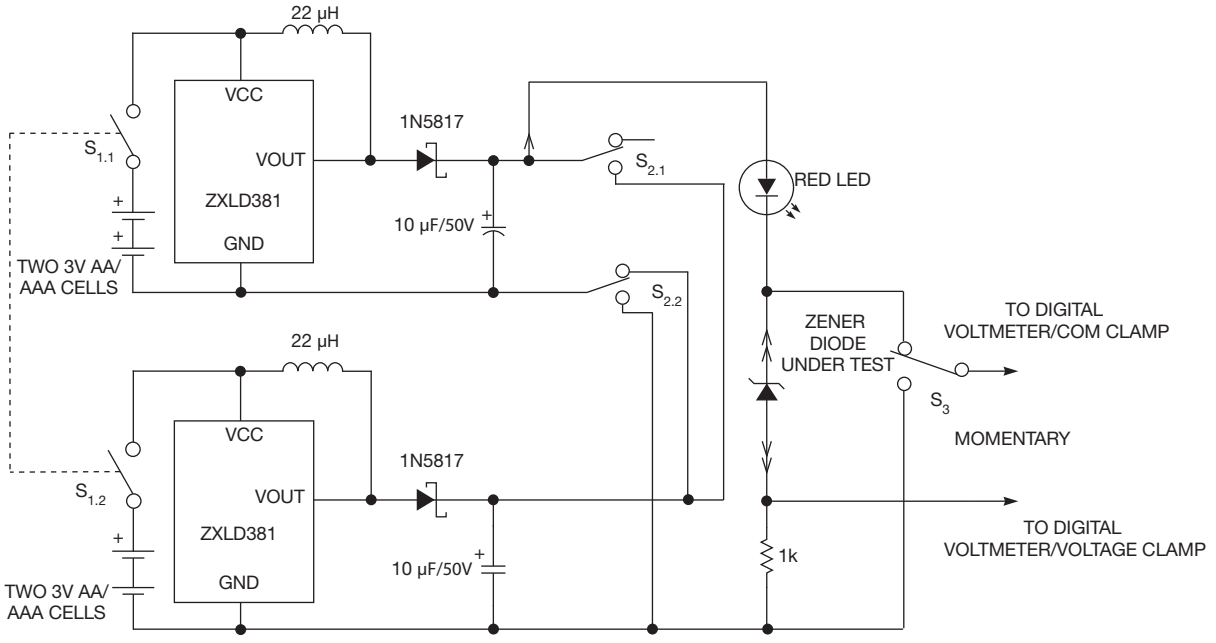

Figure 3 Two LED drivers operating in series or parallel let you boost zener voltage or current.

The red LED provides visual indication of the current flowing through the resistor-zener-diode circuit. The higher the current, the brighter the LED will light. You can omit this LED if you don't need such an indication. The voltage drop from the red LED lowers voltages by about 1.8 to 2V. If you connect the zener diode's conducting current in the forward direction, its voltage drop is equal to that of an ordinary silicon diode, or approximately 0.6 to 0.8V. When you apply forward voltage, Schottky and small-signal germanium diodes exhibit 0.2 to 0.25V and 0.35 to 0.45V drops, respectively.

Figure 4 shows an assembled zener-diode-tester circuit. The zener diode under test is On Semiconductor's (www.onsemi.com) BZX55C15RL, which has a working voltage of 13.8 to 15.6V. The voltage in **Figure 4** was measured at a zener-diode current of 2.8 mA.

You can use the circuits in **figures 2** and **3** to test LEDs, regardless of their color. Place the forward-biased LED you would like to test and set S_2 to the lower position. The voltage drop across the 1-kΩ resistor corresponds to the current in milliamps that flows through the LED under test. The circuit can light up even HB (high-brightness) LEDs because, due to their high efficacy, they start lighting at current values as low as a few milliamps. EDN

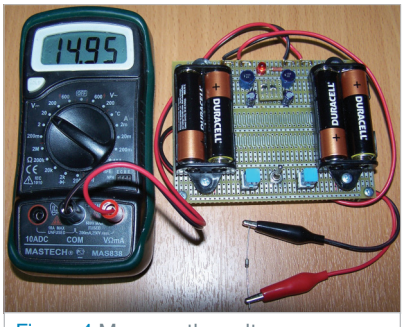

Figure 4 Measure the voltage across a zener diode with a digital multimeter.

Switched-capacitor voltage multiplier achieves 95% efficiency

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

A capacitor that you charge through a resistor operates at 50% efficiency; hence, many engineers avoid using switched-capacitor dc/dc converters. That efficiency figure holds true only for capacitors with no initial voltage, however. If you decide to switch a precharged capacitor, you can transfer energy to an output with a power-efficiency approaching 100%.

To attain a four-thirds multiple of an input supply voltage, you can charge three capacitors to one-third of each supply voltage and then add that one-third of the input voltage to the input voltage to yield an output voltage that is four-thirds of the input voltage. You series-connect three capacitors of equal value and then charge this series string to a voltage equal to the input voltage. Because the values are equal, each of the capacitors charges to one-third the input voltage. The circuit then connects these three capacitors in parallel on top of the input voltage and switches this increased voltage to the output (**Figure 1**). The circuit repeats these two phases of operation at clock frequency F.

C_{IN} and C_{OUT} are filtering capacitors at the input and output, respectively. R_P is a protective resistor, which limits the inrush current to the capacitors at power-on. As the output voltage rises, it closes IC_5 and shorts out this resistor. Schott-

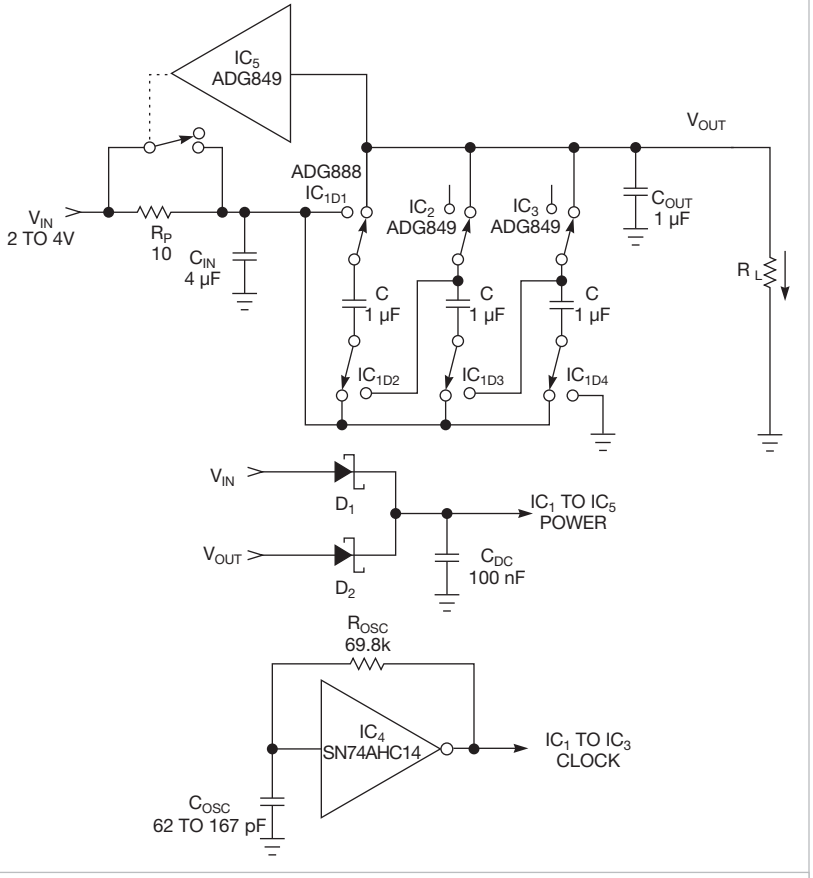

Figure 1 This charge-pump circuit raises the input voltage to a 1.33 multiple. You can use it as a 2.5-to-3.3V converter. The circuit works in environments with high magnetic fields that might interfere with an inductor-based converter.

ky diodes D_1 and D_2 allow you to power the ICs using the input voltage until the output rises, at which time the higher output voltage powers the ICs. C_{DC} is a storage and decoupling capacitor for this power bus. The higher power-supply voltage is necessary for proper operation, and it lowers the on-resistance of the analog switches. The 0.4Ω on-resistance of IC_1 results in low circuit losses and high-efficiency operation. IC_1, IC_2, and IC_3 exhibit a break-before-make operation, which is essential in this case.

For a 50%-duty-cycle clock, you can calculate the theoretical power-efficiency of the converter, according to the following **equation**:

$$\eta \simeq 1-\left(\frac{1}{2}+\frac{\frac{C_{OUT}}{3C}}{\left(1+\frac{C_{OUT}}{3C}\right)^2}\right)\times\frac{1}{12R_LCF}.$$

If the value of C_{OUT} is equal to the value of C, the power loss due to charging of the three capacitors is about two-thirds of the power loss during the discharging phase. The power consumption of the control circuit reduces the efficiency of this calculated value. For CMOS circuits, the power consumption rises linearly with the operating frequency. By choosing the operating frequency, you can optimize the efficiency of the circuit. The optimum frequency is inversely proportional to the load resistance, R_L. Fortunately, the efficiency maximum is flat, so you can achieve efficiencies higher than 90% over a wide range of values for R_L. You can attain 94% efficiency driving a 120Ω load over clock frequencies of 100 to 400 kHz. If you set a 229-kHz operating frequency, an input of 2.2V yields a 2.87V output at an efficiency of 95.9%. The optimum clock frequency shifts to lower values at lighter loads. You can drive a 269Ω load at 100 kHz and achieve an output of 2.88V. EDN

Bootstrap circuit speeds solenoid actuation

By Ralf Kelz, Seefeld, Germany

The circuit in this Design Idea bootstraps a large capacitor in series with the solenoid to provide a large actuation voltage (**Figure 1**). This higher voltage provides substantially more current to operate the solenoid (**Figure 2**, which is available at www.edn.com/100610dib), speeding the operation of the solenoid. You can also choose operating voltages or solenoid specifications that result in lower continuous current through the solenoid, reducing dc power consumption and resulting in a cooler-running solenoid with better reliability.

When there is a 0V input to the circuit, both transistors are off. Resistor R_1 slowly charges the left side of capacitor C_1 to the 24V power-supply voltage. D_2 clamps the right side of capacitor C_1 to 0.6V. When the input signal goes high, both the Q_1 and the Q_2 transistors turn on. This action quickly drives the left side of C_1 to ground. Because voltage cannot change instantaneously across a capacitor, the right side of C_1 goes down to –23.4V. D_2 steers the solenoid current into the capacitor until it discharges, at which time the solenoid current conducts through D_2 to ground. D_1 prevents a voltage-overshoot spike when the circuit turns off, and current suddenly stops flowing in D_1. It clamps the bottom leg of the solenoid to 24.6V until the current decays in the solenoid.

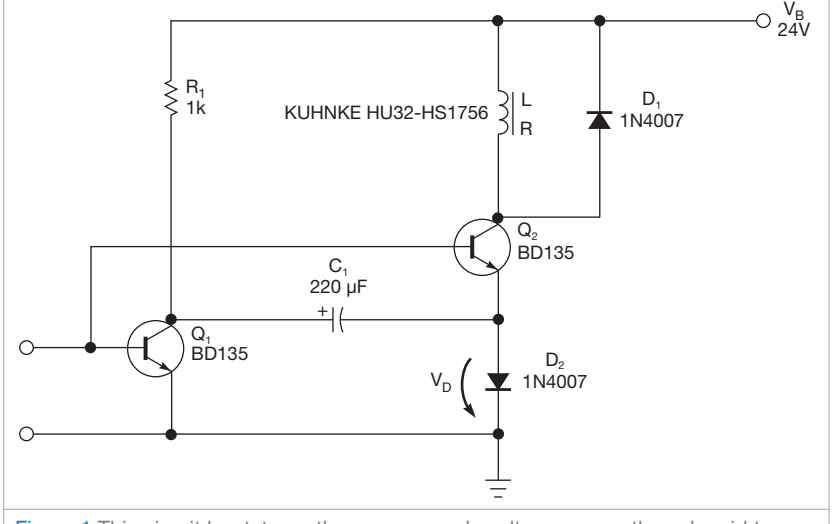

Figure 1 This circuit bootstraps the power supply voltage across the solenoid to temporarily double the actuation voltage.

THE TIME CONSTANT DEPENDS ON THE SOLENOID'S INDUCTANCE AND THE CAPACITOR'S VALUE.

The time constant of the circuit depends on the inductance of the solenoid and the value you choose for the capacitor, which you can calculate with the following **equations**:

$$I(t)=\frac{(2\,V_{IN}-V_D)\,e^{\left(-\frac{t}{\tau}\right)}\sinh(\omega t)}{\omega L};$$

$$\omega=\sqrt{\left(\frac{R}{2L}\right)^2-\frac{1}{LC}}\ ;\ \text{and}\ \tau=\frac{2L}{R}.$$

In these **equations**, e is the mathematical constant, ω is the radian angular frequency, and t is time in seconds. In addition, L is inductance and R is resistance. EDN

productroundup

POWER SOURCES

One-eighth-brick converter delivers 30A or 125W with high efficiency for 100W platforms

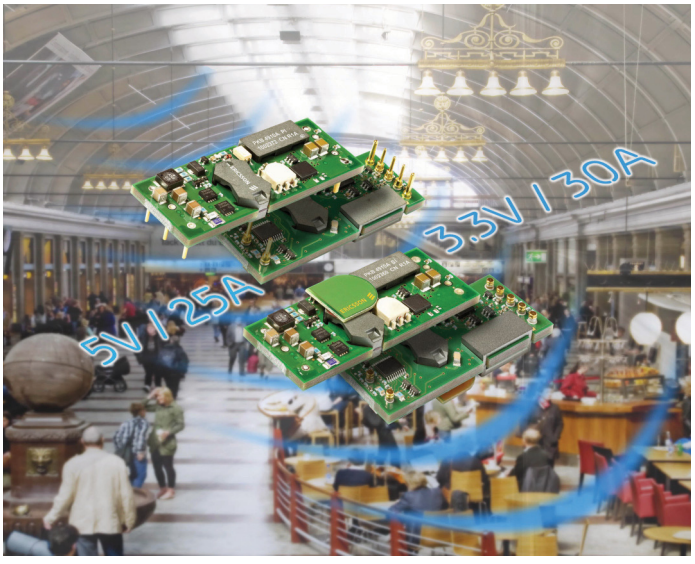

The PKB-A series of one-eighth-brick converters delivers 30A or 125W with high efficiency throughout the operating range for applications requiring an average power of 100W. The series comprises the 5V, 25A PKB 4111A and the 3.3V, 30A PKB 4910A. Both products are available for surface-mount or through-hole-assembly use. The 5V 411A has a typical efficiency of 93.1% at its nominal voltage of 48V at a 50% load and 92.3% efficiency at full load. The 3.3V 4910A offers 92.7% efficiency at a 50% load and 91% at full load. The PKB-A sells for $24 (OEM quantities).

Ericsson, www.ericsson.com

Switching ac/dc power-supply modules have 85% efficiency

Products in the CFM101S series of 100W, open-frame switching ac/dc power-supply modules measure only 2×4×1.1 in. and have universal 90 to 264V-ac inputs and output voltages of 12, 15, 20, 24, or 48V dc. The units feature power-factor correction that meets EN61000-3 specifications and achieve efficiencies as high as 85%. They also feature built-in short-circuit and overvoltage protection and convection-cooled operation in ambient temperatures of 0 to 40°C. Prices are in the low $30s (OEM quantities).

Power Sources Unlimited, www.psui.com

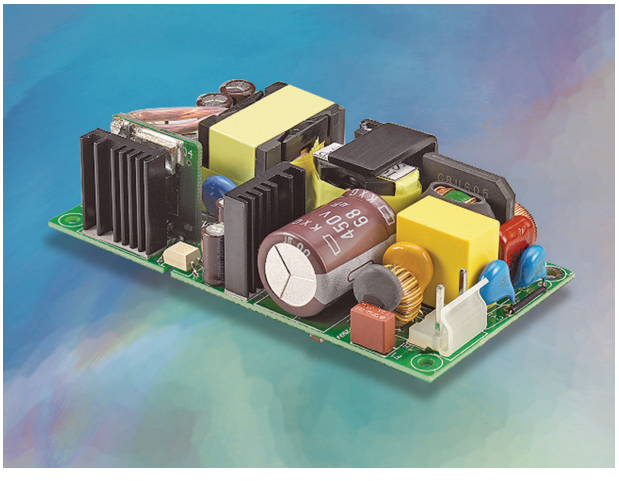

productroundup

Generic inverters target use in CCFL-backlit LCDs

The Smart Force SFW series of dc/ac inverters finds use in a range of CCFL-backlit LCD applications. The devices come in single-, dual-, and multiple-lamp configurations for use in 5.7- to 19-in.-diagonal LCDs in medical, industrial, and other applications. They also feature onboard lamp-current regulation and dimming, strike voltages as high as 2400V rms, an operating-temperature range of −20 to +85°C, and open-lamp detection. A typical distributor price for an SFWA160JF single-lamp inverter with 6-mA-rms output current and a JST SM02 (8.0) output connector is less than $10 (production quantities).

Endicott Research Group, www.ergpower.com

SENSORS/ TRANSDUCERS

Miniature infrared optoelectronic sensor has logic output

The OPL611 series of light sensors has digital outputs and comes in surface-mount packages. The sensors, which have data rates as high as 250 kbaud, comprise a Photologic chip encapsulated in transparent molded plastic, providing optical coupling and mechanical protection. The sensors measure 3×3 mm, with a height of 1 mm. They match mechanically and spectrally with the vendor's OP180, OP200, OP250, OP270, and OP280 series of infrared LEDs. Available with a buffer, a 10-kΩ pullup resistor, and open-collector outputs, the sensors feature a direct TTL/LSTTL CMOS interface. Minimum input power is 1 mW/cm^2, power-supply voltage ranges from 4.5 to 16V, and operating temperature ranges from −40 to +85°C. Typical price for the sensors is 42 cents (1000).

Optek Technology, www.optekinc.com

Digital ambient-light-sensing, proximity-detection devices operate behind darkened glass

The TSL2771 family of digital ambient-light-sensing and proximity-detection sensors eliminates the need to use clear glass or plastic in front of the sensor or drill holes in a display, bezel, or frame. These holes traditionally enable the light to reach the sensor. The TSL2771 device family instead operates behind darkened glass or other translucent materials, providing both ambient-light sensing and proximity detection. The devices incorporate an IR LED current-limited drive, analog-to-digital conversion, interrupt capability, multiple I^{2}C-interface-voltage options, and flexible programming. The TSL27711 and TSL27713 sell for $1.07 (1000). The TSL2580 with an SMBus interface and the TSL2581 with an I^{2}C interface are available in FN packages and sell for 75 cents each (1000).

Texas Advanced Optoelectronic Solutions, www.taosinc.com

CONNECTORS

PCB-terminal block offers design flexibility

The PTSM PCB-terminal block stands just 5 mm tall with a 2.5-mm center line. The spring-cage connector is rated for 6A and is available in two mounting technologies: pin-in-paste and surface mount. It targets use in LED lights, building automation, and power-supply units. The PTSM sells for 33 cents to $1.10 for connectors with two to eight positions.

Phoenix Contact,
www.phoenixcontact.com

EDN ADVERTISER INDEX

Company	Page
Agilent Technologies	6
Audio Precision	41
Austriamicrosystems AG	28
Avago	26a, 26b
Avnet Electronics Marketing	10
Avtech Electrosystems Ltd	57
Canon Communications LLC Trade Events	39
CST GmbH	C-3
Digi-Key Corp	C-1, 3
Express PCB	33
International Rectifier Corp	7
Keil Software	45
Linear Technology	42a, 42b
Maxim Intergrated Products	C-4
Memory Protection Devices	33
Mentor Graphics	13
Micrel	4
MicroPower Direct	57
Microsoft Corp	21
Mill-Max Manufacturing Corp	11
Mornsun America LLC	55
National Instruments	25
Renesas Technology Corp	C-2
RFM	49, 51, 53
Samtec USA	8
Stanford Research Systems Inc	29
Lambda Americas	27
TTI Inc	23
Vicor Corp	15, 35
XP Power	17

EDN provides this index as an additional service.The publisher assumes no liability for errors or omissions.

WARD BROWN • MICROCHIP TECHNOLOGY

An old PAL saves the day

DANIEL VASCONCELLOS

Early in my career, I worked for a defense contractor that made one-of-a-kind state-of-the-art systems. The contractor did the system design in-house, but outside contractors made many of the subsystem components per our specifications. My job on one such team was to integrate the status and control of one subsystem with the main system so that the two worked in harmony.

We lacked the luxuries that most take for granted today. Storage scopes were usually limited to two traces, and the storage was by means of a persistent screen phosphor. The scope captured and displayed slow signals better than fast signals; very fast signals sometimes required several repeated triggers to show any stored display at all.

PAL (programmable-array logic) was relatively new at the time, and tools to program these devices were crude. We could do a lot with a PAL device, but reducing the Boolean equations of the design to fit into the PAL was left to the circuit designer's wit. The designers also used a fuse map to convert the reduced equations into PAL object code to create the finished logic design. They manually entered the finished PAL code into the programmer so it could burn the fuses into the PAL. These PALs were OTP (one-time-programmable) devices, so it was a painstaking and difficult-to-verify process with the crude instruments available at the time. Much of the design effort involved being careful and hoping for the best when we had integrated everything.

When a design didn't work, we discarded the device. I remember plugging my PAL, which took weeks of effort to create, into the system and watching in disappointment as the system responded in a seemingly random fashion. How could this be? I had been so careful.

I rechecked all my equations, fuse maps, and PAL code but could find no errors. I used the crude storage scope on a few critical signals, but they seemed to be completely random. It acted as if one of the status signals was doing the opposite of what it was supposed to do. Did I invert one of the signals by mistake? I rechecked again and found that I had not. Did the vendor of the subsystem invert the signal? There was only one way to find out: I changed the PAL code to invert the signal in question. I burned another device and tried again.

This time, to my surprise and delight, the system responded flawlessly. Because this was a one-of-a-kind system and because there was no time to return the subsystem for repair, management decided to leave well enough alone and shipped the final system as it was.

About a year or so later, management called upon me to help debug a system that had come back for repair. Much had happened in the meantime, and I hardly recognized the system I was asked to help with. The problem was that the controller was acting erratically. I got out the old schematics and started to look at the interface signals. The activity looked vaguely familiar. It seemed chaotic and random.

On a hunch, I went back to my desk and looked through my junk box. I was relieved to find that what I was looking for was still there, stuck in a piece of conductive foam: my original PAL that didn't work because one status signal was inverted. I replaced the PAL in the system with that old PAL and voilà! Problem solved.

It turned out that when the customer returned the system for repairs, one of the problems was with the subsystem that my PAL interfaced with. External contractors had originally manufactured that subsystem, and the customers returned it to the vendor for repair. Someone there apparently noticed what the original people hadn't and corrected the inverted signal according to specification. EDN

Ward Brown is a staff applications engineer at Microchip Technology (Chandler, AZ).

www.edn.com/tales

Make an impact

Unleash the power of CST STUDIO SUITE®,

the No 1 technology for electromagnetic simulation.

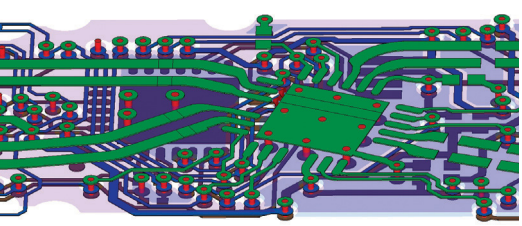

→ Get equipped with leading edge EM technology. CST's tools enable you to characterize, design and optimize electromagnetic devices all before going into the lab or measurement chamber. This can help save substantial costs especially for new or cutting edge products, and also reduce design risk and improve overall performance and profitability.

Involved in signal or power integrity analysis? You can read about how CST technology was used to simulate and optimize a digital multilayer PCB's performance at www.cst.com/pcb. If you're more interested in EMC/EMI or microwave components, we've a wide range of worked application examples live on our website at www.cst.com/apps.

Now even more choice for SI/PI simulation. The extensive range of tools integrated in CST STUDIO SUITE enables numerous applications to be analyzed without leaving the familiar CST design environment. This complete technology approach enables unprecedented simulation reliability and additional security through cross verification.

→ Grab the latest in simulation technology. Choose the accuracy and speed offered by CST STUDIO SUITE.

CST

CHANGING THE STANDARDS

CST of America®, Inc. | To request literature (508) 665 4400 | www.cst.com

USB-to-SD card bridge accelerates your design

Add Hi-Speed USB to Full-Speed applications. No software driver needed!

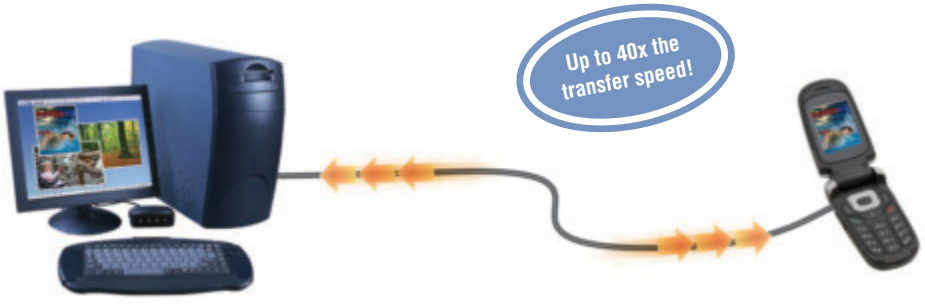

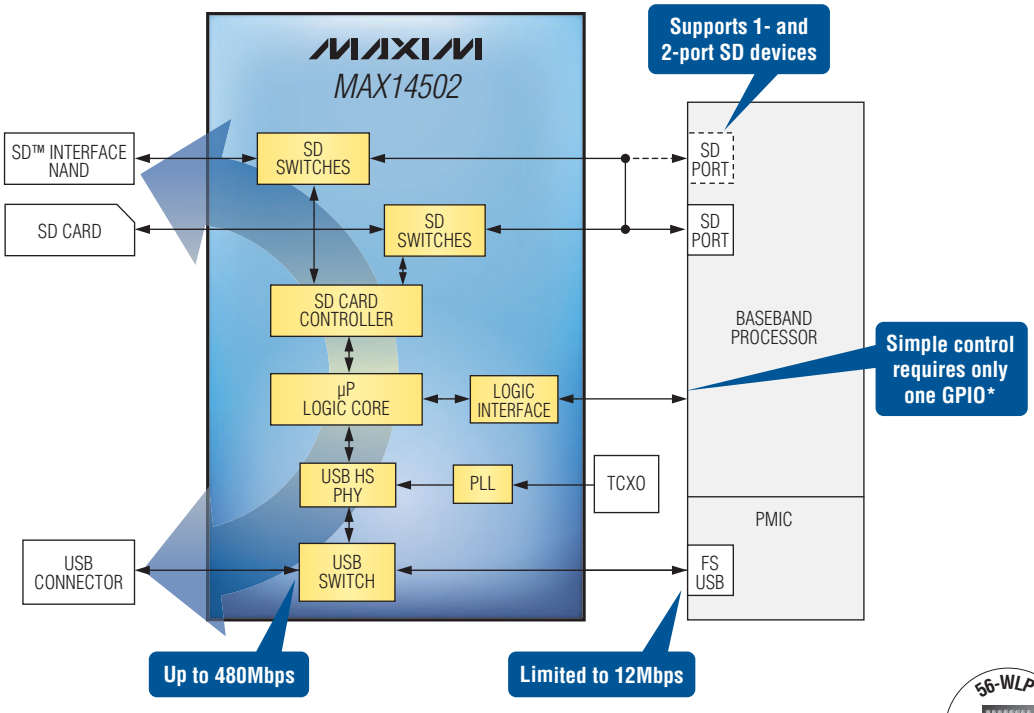

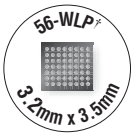

*For two-port option, control is done through an I²C interface.
†Also available in a 5mm x 5mm TQFN.
SD is a trademark of SD Card Association.

www.maxim-ic.com/MAX14502-info

MAXIM DIRECT.
www.maxim-ic.com/shop

AVNET electronics marketing
www.em.avnet.com/maxim

MAXIM
INNOVATION DELIVERED™

For free samples or technical support, visit our website.